

MBZUAI

Digital.Commons@MBZUAI

Computer Vision Faculty Publications

Scholarly Works

1-24-2022

Transformers in Medical Imaging: A Survey

Fahad Shamshad

Mohamed Bin Zayed University of Artificial Intelligence

Salman Khan

Mohamed Bin Zayed University of Artificial Intelligence

Syed Waqas Zamir

The Inception Institute of Artificial Intelligence

Muhammad Haris Khan

Mohamed Bin Zayed University of Artificial Intelligence

Munawar Hayat

Monash University

See next page for additional authors

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/cvfp>



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Biomedical Engineering and Bioengineering Commons](#)

Preprint: arXiv

- Archived with thanks to arXiv
- Preprint License: [CC by SA](#)
- Uploaded 24 March 2022

Recommended Citation

F. Shamshad, et.al., "Transformers in medical imaging: a survey", 2022, arXiv:2201.09873

This Article is brought to you for free and open access by the Scholarly Works at Digital.Commons@MBZUAI. It has been accepted for inclusion in Computer Vision Faculty Publications by an authorized administrator of Digital.Commons@MBZUAI. For more information, please contact libraryservices@mbzuai.ac.ae.

Authors

Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan, Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu

Transformers in Medical Imaging: A Survey

Fahad Shamshad, Salman Khan, Syed Waqas Zamir, Muhammad Haris Khan,
Munawar Hayat, Fahad Shahbaz Khan, and Huazhu Fu

Abstract—Following unprecedented success on the natural language tasks, Transformers have been successfully applied to several computer vision problems, achieving state-of-the-art results and prompting researchers to reconsider the supremacy of convolutional neural networks (CNNs) as de facto operators. Capitalizing on these advances in computer vision, the medical imaging field has also witnessed growing interest for Transformers that can capture global context compared to CNNs with local receptive fields. Inspired from this transition, in this survey, we attempt to provide a comprehensive review of the applications of Transformers in medical imaging covering various aspects, ranging from recently proposed architectural designs to unsolved issues. Specifically, we survey the use of Transformers in medical image segmentation, detection, classification, reconstruction, synthesis, registration, clinical report generation, and other tasks. In particular, for each of these applications, we develop taxonomy, identify application-specific challenges as well as provide insights to solve them, and highlight recent trends. Further, we provide a critical discussion of the field's current state as a whole, including the identification of key challenges, open problems, and outlining promising future directions. We hope this survey will ignite further interest in the community and provide researchers with an up-to-date reference regarding applications of Transformer models in medical imaging. Finally, to cope with the rapid development in this field, we intend to regularly update the relevant latest papers and their open-source implementations at <https://github.com/fahadshamshad/awesome-transformers-in-medical-imaging>.

Index Terms—Transformers, medical image analysis, vision transformers, deep neural networks, clinical report generation

1 INTRODUCTION

CONVOLUTIONAL Neural Networks (CNNs) [1]–[4] have significantly impacted the field of medical imaging due to their ability to learn highly complex representations in a data-driven manner. Since their renaissance, CNNs have demonstrated remarkable improvements on numerous medical imaging modalities, including Radiography [5], Endoscopy [6], Computed Tomography (CT) [7], [8], Mammography Images (MG) [9], Ultrasound Images [10], Magnetic Resonance Imaging (MRI) [11], [12], and Positron Emission Tomography (PET) [13], to name a few. The workhorse in CNNs is the *convolution* operator, which operates locally and provides translational equivariance. While these properties help in developing efficient and generalizable medical imaging solutions, the local receptive field in convolution operation limits capturing long-range pixel relationships. Furthermore, the convolutional filters have stationary weights that are not adapted for the given input image content at inference time.

Meanwhile, significant research effort has been made by the vision community to integrate the attention mechanisms [14]–[16] in CNN-inspired architectures [17]–[22]. These attention-based ‘Transformer’ models have become an attractive solution due to their ability to encode long-range dependencies and learn highly effective feature representa-

tions [23]. Recent works have shown that these Transformer modules can fully replace the standard convolutions in deep neural networks by operating on a sequence of image patches, giving rise to Vision Transformers (ViTs) [22]. Since their inception, ViT models have been shown to push the state-of-the-art in numerous vision tasks, including image classification [22], object detection [24], semantic segmentation [25], image colorization [26], low-level vision [27], and video understanding [28] to name a few. Furthermore, recent research indicate that the prediction errors of ViTs are more consistent with those of humans than CNNs [29]–[32]. These desirable properties of ViTs have sparked great interest in the medical community to adapt them for medical imaging applications, thereby mitigating the inherent inductive biases of CNNs [33].

Motivation and Contributions: Recently, medical imaging community has witnessed an exponential growth in the number of Transformer based techniques, especially after the inception of ViTs (see Fig. 1). The topic is now dominant in prestigious medical imaging conferences and journals, and it is getting increasingly difficult to keep pace with the recent progress due to the rapid influx of papers. As such, a survey of the existing relevant works is timely to provide a comprehensive account of new methods in this emerging field. To this end, we provide a holistic overview of the applications of Transformer models in medical imaging. We hope this work can provide a roadmap for the researchers to explore the field further. Our major contributions include:

- This is the first survey paper that comprehensively covers applications of Transformers in the medical imaging domain, thereby bridging the gap between the vision and medical imaging community in this rapidly evolving area. Specifically, we present a comprehensive overview of more than 125 relevant

F. Shamshad, S. Khan, M. H. Khan and F. S. Khan are with the MBZ University of Artificial Intelligence, Abu Dhabi, UAE.

E-mail: firstname.lastname@mbzuai.ac.ae

S. W. Zamir is with the Inception Institute of Artificial Intelligence, Abu Dhabi, UAE.

M. Hayat is with the Faculty of IT, Monash University, Clayton VIC 3800, Australia.

S. Khan is also with the CECS, Australian National University, Canberra ACT 0200, Australia.

F. S. Khan is also with the Computer Vision Laboratory, Linköping University, Sweden.

H. Fu is with the Institute of High Performance Computing, Agency for Science, Technology and Research (A*STAR), Singapore.




Figure 1: (Left) The pie-charts show statistics of the papers included in this survey according to medical imaging problem settings and data modalities. The rightmost figure shows consistent growth in the recent literature (for year 2021). Seg: segmentation, Class: classification, Recons: reconstruction, Reg: registration, Synth: synthesis, Det: detection, Rep: report generation, US: ultrasound.




Figure 2: A diverse set of application areas of Transformers in medical imaging covered in this survey.

papers to cover the recent progress.

- We provide a detailed coverage of the field by categorizing the papers based on their applications in medical imaging as depicted in Fig. 2. For each of these applications, we develop a taxonomy, highlight task-specific challenges, and provide insights about solving them based on the literature reviewed.
- Finally, we provide a critical discussion of the field’s current state as a whole, including identifying key challenges, highlighting open problems, and outlining promising future directions.
- Although the main focus of this survey is on Vision Transformers, we are also the first since the inception of the original Transformer, about half a decade ago, to extensively cover its language modeling capabilities in the clinical report generation task (see Sec. 9).

Paper Organization. The rest of the paper is organized as follows. In Sec. 2, we provide background of the field with a focus on salient concepts underlying Transformers. From

Sec. 3 to Sec. 10, we comprehensively cover applications of Transformers in several medical imaging tasks as shown in Fig. 2. In particular, for each of these tasks, we develop a taxonomy and identify task-specific challenges. Sec. 11 presents open problems and future directions about the field as a whole. Finally, in Sec. 12, we give recommendations to cope with the rapid development of the field and conclude the paper.

2 BACKGROUND

Medical imaging approaches have undergone significant advances over the past few decades. In this section, we briefly provide a background of these advancements and broadly group them into three categories: hand-crafted, CNN-based, and ViT-based. For the hand-crafted and CNN-based methods, we describe the underlying working principles along with their major strengths and shortcomings in the context of medical imaging. For the ViT-based methods, we highlight core concepts behind their success and defer further details to later sections.

2.1 Hand-Crafted Approaches

Conventional algorithms to solve medical imaging tasks are based on hand-crafted mathematical models designed by field experts using domain knowledge. The development of these hand-crafted models with a focus on refining discriminative features and efficient optimization algorithms for a range of medical imaging problems has been the central research topic in the past [40], [41]. Successful hand-crafted models in medical imaging include total-variation [42], non-local self-similarity [43], sparsity/structured sparsity [44], Markov-tree models on wavelet coefficients [45], and untrained neural networks [46]–[48]. These models have been extensively leveraged in medical domain for image segmentation [49], reconstruction [50], disease classification [51], enhancement [52], and anomaly detection [53] due to their interpretability with solid mathematical foundations and theoretical supports on the robustness, recovery, and complexity [54], [55]. Further, unlike deep learning-based approaches, they do not require large annotated medical imaging datasets for training. This reduced reliance on labeled datasets is crucial to the medical research community

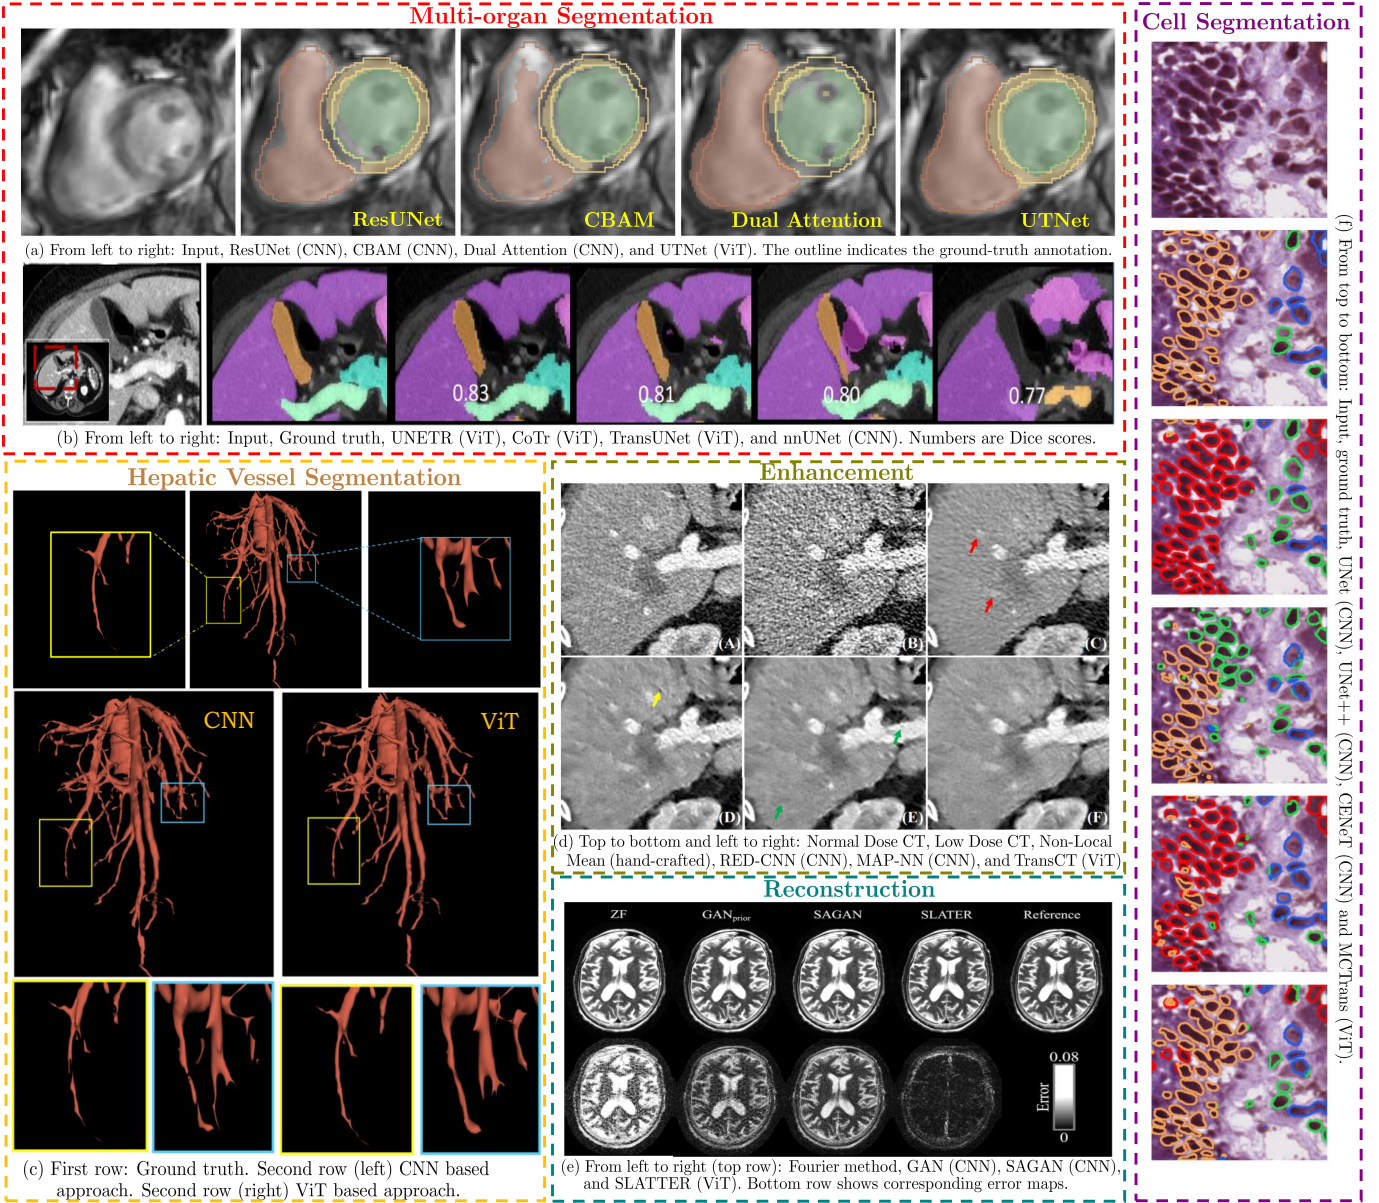


Figure 3: Applications of ViTs in various medical imaging problems along with the baseline CNN-based approaches. ViT-based approaches give superior performance as compared to CNN-based methods due to their ability to model the global context. Figure sources: (a) [34], (b) [35], (c) [36], (d) [37], (e) [38], (f) [39].

as collecting voluminous, reliable, and labeled data in the medical domain is difficult due to the lack of expert annotators, high time consumption, ethical considerations, and financial costs. However, due to inadequacy to leverage the expressive power of large medical imaging data sets, these hand-crafted models often suffer from poor discriminative capability [56]. Consequently, these models often fail to represent nuances of high-dimensional complex medical imaging data that can hamper the performance of the medical imaging diagnosis systems [57], [58]. To circumvent the poor discriminability and generalization issue, learned hand-crafted models have been proposed to exploit data better. The representative approaches include optimal directions [59], K-SVD [60], data-driven tight frame [61], low-rank models [62], and piece-wise smooth image model [63]. Next, we

explain the popular data-driven approaches explored in the literature.

2.2 CNN-based methods

CNNs are effective at learning discriminative features and extracting generalizable priors from large-scale medical datasets, thus providing excellent performance on medical imaging tasks, making them an integral component of modern AI-based medical imaging systems. The advancements in CNNs have been mainly fueled by novel architectural designs, better optimization procedures, availability of special hardware (e.g., GPUs) and purpose-built open source software libraries [64]–[66]. We refer interested readers to comprehensive survey papers related to CNNs applications in medical imaging [56], [67]–[75]. Despite considerable perfor-

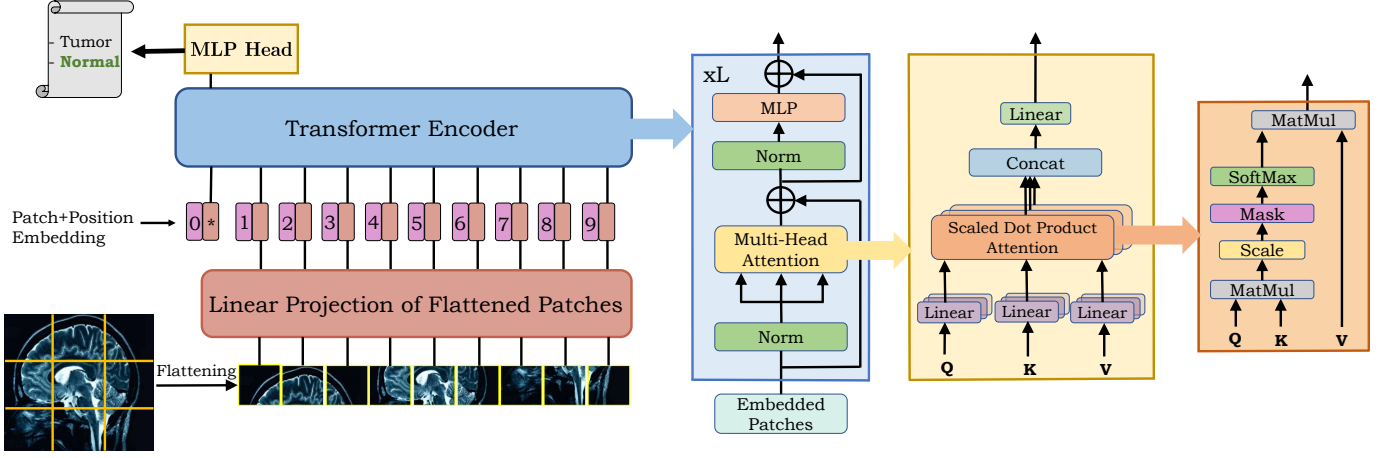


Figure 4: Architecture of the Vision Transformer (*on the left*) and details of the Vision Transformer encoder block (*on the right*). Vision transformer first splits the input image into patches and projects them (after flattening) into a feature space where a transformer encoder processes them to produce the final classification output.

mance gains, the reliance of CNNs on large labeled datasets limits their applicability over the full spectrum of medical imaging tasks. Furthermore, CNNs-based approaches are generally more challenging to interpret and often act as black box solutions. Therefore, there has been an increasing effort in the medical imaging community to amalgamate the strengths of hand-crafted and CNNs based methods resulting in the prior information-guided CNNs models [76]. These hybrid methods contain special domain-specific layers, and include unrolled optimization [77], generative models [78], and learned denoiser-based approaches [79]. Despite these architectural and algorithmic advancements, the decisive factor behind CNNs success has been primarily attributed to their image-specific inductive bias in dealing with scale invariance and modeling local visual structures. While this intrinsic locality (limited receptive field) brings efficiency to CNNs, it impairs their ability to capture long-range spatial dependencies in an input image, thereby stagnating performance [33] (see Fig. 3). This demands an alternative architectural design capable of modeling long-range pixel relationships for better representation learning.

2.3 Transformers

Transformers were introduced by Vaswani *et al.* [14] as a new attention-driven building block for machine translation. Specifically, these attention blocks are neural network layers that aggregate information from the entire input sequence [80]. Since their inception, these models have demonstrated state-of-the-art performance on several Natural Language Processing (NLP) tasks, thereby becoming the default choice over recurrent models. In this section, we will focus on Vision Transformers (ViTs) [22] that are built on vanilla Transformer model [14] by cascading multiple transformer layers to capture the global context of an input image. Specifically, Dosovitskiy *et al.* [22] interpret an image as a sequence of patches and process it by a standard transformer encoder as used in NLP. These ViT models continue the long-lasting trend of removing hand-crafted visual features and inductive biases from models in an effort to leverage the availability of larger datasets coupled with increased computational

Algorithm 1 ViT Working Principle

- 1: Split a medical image into patches of fixed sizes
- 2: Vectorize image patches via flattening operation
- 3: Create lower-dimensional linear embedding from vectorized patches via trainable linear layer
- 4: Add positional encoding to lower dimensional linear embeddings
- 5: Feed the sequence to ViT encoder as shown in Figure 4
- 6: Pre-train the ViT model on a large-scale image dataset
- 7: Fine-tune on the down stream medical image classification task

capacity. ViTs have garnered immense interest in the medical imaging community, and a number of recent approaches have been proposed which build upon ViTs. We highlight the working principle of ViT in a step-by-step manner in Algorithm 1 for medical image classification.

Below, we briefly describe the core components behind the success of ViTs that are *self-attention* and *multi-head self-attention*. For a more in-depth analysis of numerous ViT architectures and applications, we refer interesting readers to the recent relevant survey papers [23], [81]–[84].

2.3.1 Self-Attention

The success of the Transformer models has been widely attributed to the self-attention (SA) mechanism due to its ability to model long-range dependencies. The key idea behind the SA mechanism is to learn self-alignment, that is, to determine the relative importance of a single token (patch embedding) with respect to all other tokens in the sequence [80]. For 2D images, we first reshape the image $\mathbf{x} \in \mathbb{R}^{H \times W \times C}$ into a sequence of flattened 2D patches $\mathbf{x}_p \in \mathbb{R}^{N \times (P^2 C)}$, where H and W denotes height and width of the original image respectively, C is the number of channels, $P \times P$ is the resolution of each image patch, and $N = HW/P^2$ is the resulting number of patches. These flattened patches are projected to D dimension via trainable linear projection layer and can be represented in matrix form as $\mathbf{X} \in \mathbb{R}^{N \times D}$. The goal of self-attention is to capture the interaction amongst all

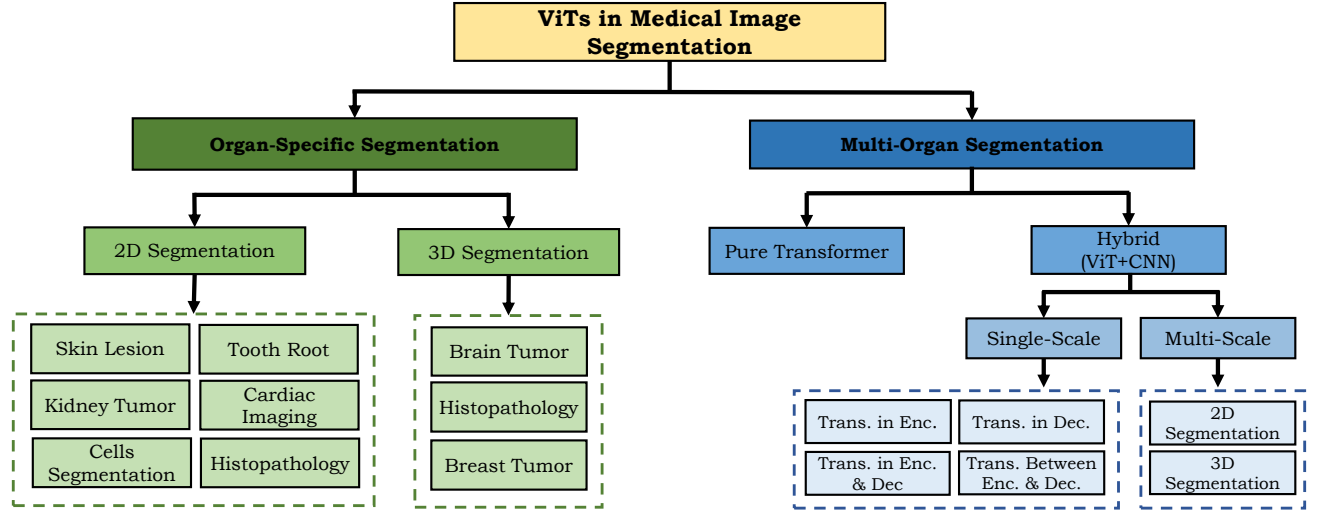


Figure 5: Taxonomy of ViT-based medical image segmentation approaches.

these N embeddings, that is done by defining three learnable weight matrices to transform input \mathbf{X} into Queries (via $\mathbf{W}^Q \in \mathbb{R}^{D \times D_q}$), Keys (via $\mathbf{W}^K \in \mathbb{R}^{D \times D_k}$) and Values (via $\mathbf{W}^V \in \mathbb{R}^{D \times D_v}$), where $D_q = D_k$. The input sequence \mathbf{X} is first projected onto these weight matrices to get $\mathbf{Q} = \mathbf{XW}^Q$, $\mathbf{K} = \mathbf{XW}^K$ and $\mathbf{V} = \mathbf{XW}^V$. The corresponding attention matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ can be written as,

$$\mathbf{A} = \text{softmax} \left(\frac{\mathbf{QK}^T}{\sqrt{D_q}} \right).$$

The output $\mathbf{Z} \in \mathbb{R}^{N \times D_v}$ of the SA layer is then given by,

$$\mathbf{Z} = \mathbf{SA}(\mathbf{X}) = \mathbf{AV}.$$

2.3.2 Multi-Head Self-Attention

Multi-Head Self Attention (MHSA) consists of multiple SA blocks (heads) concatenated together channel-wise to model complex dependencies between different elements in the input sequence. Each head has its own learnable weight matrices denoted by $\{\mathbf{W}^{Q_i}, \mathbf{W}^{K_i}, \mathbf{W}^{V_i}\}$, where $i = 0 \dots (h-1)$ and h denotes total number of heads in MHSA block. Specifically, we can write,

$$\text{MHSA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = [\mathbf{Z}_0, \mathbf{Z}_1, \dots, \mathbf{Z}_{h-1}] \mathbf{W}^O,$$

whereas $\mathbf{W}^O \in \mathbb{R}^{h \cdot D_v \times N}$ computes linear transformation of heads and \mathbf{Z}_i can be written as,

$$\mathbf{Z}_i = \text{softmax} \left(\frac{\mathbf{QW}^{Q_i} (\mathbf{KW}^{K_i})^T}{\sqrt{D_q/h}} \right) \mathbf{VW}^{V_i}.$$

Note that the complexity of computing the softmax for SA block is quadratic with respect to the length of the input sequence that can limit its applicability to high-resolution medical images. Recently, numerous efforts have been made to reduce complexity, including sparse attention [85], linearization attention [86], low-rank attention [87], memory compression based approaches [88], and improved MHSA [89]. We will discuss the efficient SA in the context of medical imaging in the relevant sections.

Further, we find it important to clarify that several alternate attention approaches [90]–[93] have been explored

in the literature based on convolutional architectures. In this survey, we focus on the specific attention used in transformer blocks (MHSA) which has recently gained significant research attention in medical image analysis. Next, we outline these methods categorized according to specific application domains.

3 MEDICAL IMAGE SEGMENTATION

Accurate medical image segmentation is a crucial step in computer-aided diagnosis, image-guided surgery, and treatment planning. The global context modeling capability of Transformers is crucial for accurate medical image segmentation because the organs spread over a large receptive field can be effectively encoded by modeling the relationships between spatially distant pixels (e.g., lungs segmentation). Furthermore, the background in medical scans is generally scattered (e.g., in ultrasound scan [94]); therefore, learning global context between the pixels corresponding to the background can help the model in preventing misclassification.

Below, we highlight various attempts to integrate ViT-based models for medical image segmentation. We broadly classify the ViT-based segmentation approaches into *organ-specific* and *multi-organ* categories, as depicted in Fig. 5, due to the varying levels of context modeling required in both sets of methods.

3.1 Organ-Specific Segmentation

ViT-based organ-specific approaches generally consider a specific aspect of the underlying organ to design architectural components or loss functions. We mention specific examples of such design choices in this section. We have further categorized organ-specific categories into 2D and 3D-based approaches depending on the input type.

3.1.1 2D Segmentation

Here, we describe the organ-specific ViT-based segmentation approaches for 2D medical scans.

Skin Lesion Segmentation. Accurate skin lesion segmentation for identifying melanoma (cancer cells) is crucial

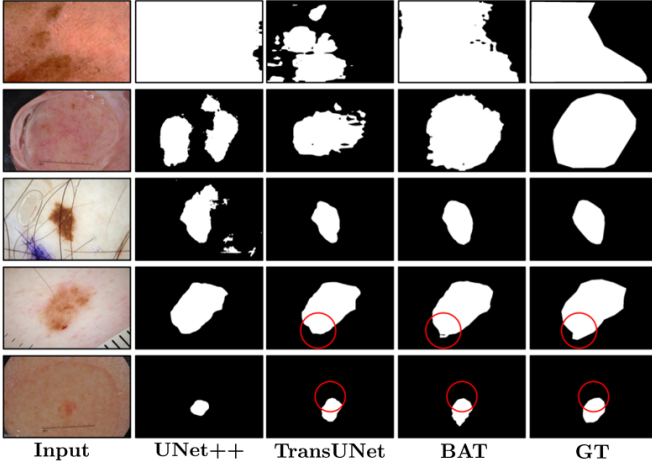


Figure 6: Comparison of different skin lesion segmentation approaches. From left to right: Input image, CNN based UNet++ [95], ViT-based TransUNet [96], Boundary aware transformer (BAT) [97], and the ground truth (GT) image. Red circles highlight small regions with an ambiguous boundary where BAT can perform well due to the use of boundary-wise prior knowledge. Image taken from [97].

for cancer diagnosis and subsequent treatment planning. However, it remains a challenging task due to significant variations in color, size, occlusions, and contrast of skin lesion areas, resulting in ambiguous boundaries [98] and consequently deterioration in segmentation performance. To address the issue of ambiguous boundaries, Wang *et al.* [97] propose a novel Boundary-Aware Transformer (BAT). Specifically, they design a boundary-wise attention gate in Transformer architecture to exploit the prior knowledge about boundaries. The auxiliary supervision of the boundary-wise attention gate provides feedback to train BAT effectively. Extensive experiments on ISIC 2016+PH2 [99], [100] and ISIC 2018 [101] validate the efficacy of their boundary-wise prior, as shown in Fig. 6. Similarly, Wu *et al.* [102] propose a dual encoder-based feature adaptive transformer network (FAT-Net) that consists of CNN and transformer branches in the encoder. To effectively fuse the features from these two branches, a memory-efficient decoder and feature adaptation module have been designed. Experiments on ISIC 2016-2018 [99], [101], [103], and PH2 [100] datasets demonstrate the effectiveness of FAT-Net fusion modules.

Tooth Root Segmentation. Tooth root segmentation is one of the critical steps in root canal therapy to treat periodontitis (gum infection) [104]. However, it is challenging due to blurry boundaries and overexposed and underexposed images. To address these challenges, Li *et al.* [105] propose Group Transformer U-Net (GT U-Net) that consists of transformer and convolutional layers to encode global and local context, respectively. A shape-sensitive Fourier Descriptor loss function [106] has been proposed to deal with the fuzzy tooth boundaries. Furthermore, grouping and bottleneck structure has been introduced in the GT U-Net to significantly reduce the computational cost. Experiments on their in-house Tooth Root segmentation dataset with six evaluation metrics demonstrate the effectiveness of GT U-Net architectural components and Fourier-based loss

function. In another work, Li *et al.* [107] propose anatomy-guided multibranch Transformer (AGMB-Transformer) to incorporate the strengths of group convolutions [108] and progressive Transformer network. Experiments on their self-collected dataset of 245 tooth root X-ray images show the effectiveness of AGMB-Transformer.

Cardiac Image Segmentation. Despite their impressive performance in medical image segmentation, Transformers are computationally demanding to train and come with a high parameter budget. To handle these challenges for cardiac image segmentation task, Deng *et al.* [109] propose TransBridge, a lightweight parameter-efficient hybrid model. TransBridge consists of Transformers and CNNs based encoder-decoder structure for left ventricle segmentation in echocardiography. Specifically, the patch embedding layer of the Transformer has been re-designed using the shuffling layer [110] and group convolutions to significantly reduce the number of parameters. Extensive experiments on large-scale left ventricle segmentation dataset, echo-cardiographs [111] demonstrate the benefit of TransBridge over CNNs and Transformer-based baseline approaches [112].

Kidney Tumor Segmentation. Accurate segmentation of kidney tumors via computer diagnosis systems can reduce the effort of radiologists and is a critical step in related surgical procedures. However, it is challenging due to varying kidney tumor sizes and the contrast between tumors and their anatomical surroundings. To address these challenges, Shen *et al.* [117] propose a hybrid encoder-decoder architecture, COTR-Net, that consists of convolution and transformer layers for end-to-end kidney, kidney cyst, and kidney tumor segmentation. Specifically, the encoder of COTR-Net consists of several convolution-transformer blocks, and the decoder comprises several up-sampling layers with skip connections from the encoder. The encoder weights have been initialized using a pre-trained ResNet [118] architecture to accelerate convergence, and deep supervision has been exploited in the decoder layers to boost segmentation performance. Furthermore, the segmentation masks are refined using morphological operations as a post-processing step. Extensive experiments on the Kidney Tumor Segmentation dataset (KiTS21) [119] demonstrate the effectiveness COTR-Net.

Cell Segmentation. Inspired from the Detection Transformers (DETR) [132], Zhang *et al.* proposed Cell-DETR [133], a Transformer-based framework for instance segmentation of biological cells. Specifically, they integrate a dedicated attention branch to the DETR framework to obtain instance-wise segmentation masks in addition to box predictions. During training, focal loss [134] and Sorenson dice loss [132] are used for the segmentation branch. To enhance performance, they integrate three residual decoder blocks [118] in Cell-DETR to generate accurate instance masks. Experiments on their in-house yeast cells dataset demonstrate the effectiveness of Cell-DETR relative to U-Net based baselines [114]. Similarly, existing medical imaging segmentation approaches generally struggle for Corneal endothelial cells due to blurry edges caused by the subject’s movement [135]. This demands preserving more local details and making full use of the global context. Considering these attributes, Zhang *et al.* [136] propose a Multi-Branch hybrid Transformer Network (MBT-Net) consisting of convolutional




Figure 7: Segmentation results for hyperspectral pathology dataset using Spectral Transformer (SpecTr). From left to right: Input image, Ground truth label, HIS Hyper (CNN-based) [113], UNet (CNN-based) [114], Attn UNet (CNN-based) [115], UNet++ (CNN-based) [95], and SpecTr (ViT-based) [116]. Image adapted from [116].

CNN-based Models	40x Magnification	20x Magnification
PSPNet [120]	0.58±0.33	0.49±0.27
U-Net [114]	0.65±0.24	0.60±0.31
DeepLabV3 [121]	0.63±0.28	0.67±0.24
FPN [122]	0.64±0.20	0.72±0.22
PAN [123]	0.63±0.24	0.69±0.23
LinkNet [124]	0.35±0.33	0.54±0.25
Transformer-based Models	40x Magnification	20x Magnification
TransUNet [96]	0.77±0.12	0.77±0.13
Swin-UNet [125]	0.53±0.23	0.42±0.23
Swin Transformer (Base) [126]	0.79±0.14	0.71±0.26
Segmenter [127]	0.80±0.14	0.82±0.11
Medical Transformer [128]	0.71±0.14	0.62±0.17
BEiT [129]	0.72±0.21	0.66±0.28

Table 1: Evaluation of Transformer-based semantic segmentation methods for pathological image segmentation in terms of average jaccard index on PAIP liver histopathological dataset [130]. It can be seen that transformer-based models tend to outperform CNNs with the exception of Swin-UNet. Results are from [131], which is one of the first study to systematically evaluate the performance of transformers on pathological image segmentation task.

and transformer layers. Specifically, they propose a body-edge branch that provides precise edge location information and promotes local consistency. Extensive ablation studies on their self-collected TM-EM3000 and public Alisarine dataset [137] of Corneal Endothelial Cells show the effectiveness of MBT-Net architectural components.

Histopathology. Histopathology refers to the diagnosis and study of the diseases of tissues under a microscope and is the gold standard for cancer recognition. Therefore accurate automatic segmentation of histopathology images can substantially alleviate the workload of pathologists. Recently, Nguyen *et al.* [131] systematically evaluate the performance of six latest ViTs, and CNNs-based approaches on whole slide images of the PAIP liver histopathological dataset [130]. Their results (shown in Table 1) demonstrate that almost all Transformer-based models indeed exhibit superior performance as compared to CNN-based approaches due to their ability to encode the global context.

3.1.2 3D Medical Segmentation

Here, we describe ViT-based segmentation approaches for volumetric medical data.

Brain Tumor Segmentation. An automatic and accurate brain tumor segmentation approach can lead to the timely diagnosis of neurological disorders such as Alzheimer’s disease. Recently, ViT-based models have been proposed to segment brain tumors effectively. Wang *et al.* [138] have made

the first attempt to leverage Transformers for 3D multimodal brain tumor segmentation by effectively modeling local and global features in both spatial and depth dimensions. Specifically, their encoder-decoder architecture, TransBTS, employs a 3D CNN to extract local 3D volumetric spatial features and Transformers to encode global features. Progressive upsampling in the 3D CNN-based decoder has been used to predict the final segmentation map. To further boost the performance, they make use of test-time augmentation. Extensive experimentation on BraTS 2019¹ and BraTS 2020² datasets show the effectiveness of their proposed approach compared to CNN-based methods. Unlike most of the ViT-based image segmentation approaches, TransBTS does not require pre-training on large datasets and has been trained from scratch. In another work, inspired from the architectural design of TransBTS [138], Jia *et al.* [139] propose Bi-Transformer U-Net (BiTr-UNet) that performs relatively better on BraTS 2021 [140] segmentation challenge. Different from TransBTS [138], BiTr-UNet consists of an attention module to refine encoder and decoder features and has two ViT layers (instead of one as in TransBTS). Furthermore, BiTr-UNet adopts a post-processing strategy to eliminate a volume of predicted segmentation if the volume is smaller than a threshold [141] followed by model ensemble via majority voting [142]. Similarly, Peiris *et al.* [143] propose a light-weight UNet shaped volumetric transformer, VT-UNet, to segment 3D medical image modalities in a hierarchical manner. Specifically, two self-attention layers have been introduced in the encoder of VT-UNet to capture both global and local contexts. Furthermore, the introduction of window-based self-attention and cross-attention modules and Fourier positional encoding in the decoder significantly improve the accuracy and efficiency of VT-UNet. Experiments on BraTs 2021 [140] show that VT-UNet is robust to data artifacts and exhibits strong generalization ability. In another similar work, Hatamizadeh *et al.* [145] propose Swin UNet based architecture, Swin UNETR, that consists of Swin transformer as the encoder and a CNN-based decoder. Specifically, Swin UNETR computes self-attention in an efficient shifted window partitioning scheme and is a top-performing model on BraTs 2021 [140] validation set. In Table 2, we provide dice score and other parameters of various Transformer based models for the 3D multimodal BraTs 2021 dataset [140].

Histopathology. Boxiang *et al.* [116] propose Spectral Transformer (SpecTr) for hyperspectral pathology image

1. <https://www.med.upenn.edu/cbica/brats2019>

2. <https://www.med.upenn.edu/cbica/brats2020>

Method	#params	Flops	Dice Score (Avg.)
TransBTS [138]	33 M	333 G	84.99
BiTr-UNet [139]	-	-	86.20
UNETR [35]	102.5 M	193.5 G	84.51
nnFormer [144]	39.7 M	110.7 G	86.56
Swin UNETR [145]	61.98 M	394.84 G	88.97
VT-UNET-T [143]	5.4 M	52 G	86.82
VT-UNET-S [143]	11.8 M	100.8 G	87.00
VT-UNET-B [143]	20.8 M	165 G	88.07

Table 2: Segmentation results and parameters of various Transformer-based models on 3D Multimodal Brain Tumor BraTS 2021 dataset [140].

segmentation, which employs transformers to learn the contextual feature across the spectral dimension. To discard the irrelevant spectral bands, they introduce a sparsity-based scheme [146]. Furthermore, they employ separate group normalization for each band to eliminate the interference caused by distribution mismatch among spectral images. Extensive experimentation on the hyperspectral pathology dataset, Cholangiocarcinoma [147], shows the effectiveness of SpecTr as also shown in Fig. 7.

Breast Tumor Segmentation. Detection of breast cancer in the early stages can reduce the fatality rate by more than 40% [148]. Therefore, automatic breast tumor detection is of immense importance to doctors. Recently, Zhu *et al.* [149] propose a region aware transformer network (RAT-Net) to effectively fuse the Breast tumor region information into multiple scales to obtain precise segmentation. Extensive experiments on a large ultrasound breast tumor segmentation dataset show that RAT-Net outperforms CNN and transformer-based baselines. Similarly, Liu *et al.* [150] also propose a hybrid architecture consisting of transformer layers in the decoder part of 3D UNet [151] to effectively segment tumors from volumetric breast data.

3.2 Multi-organ Segmentation

Multi-organ segmentation aims to segment several organs simultaneously and is challenging due to inter-class imbalance and varying sizes, shapes, and contrast of different organs. ViT models are particularly suitable for the multi-organ segmentation due to their ability to effectively model global relations and differentiate multiple organs. We have categorized multi-organ segmentation approaches based on the architectural design, as these approaches do not consider any organ-specific aspect and generally focus on boosting performance by designing effective and efficient architectural modules [152]. We categorize multi-organ segmentation approaches into *Pure Transformer* (only ViT layers) and *Hybrid Architectures* (both CNNs and ViTs layers).

3.2.1 Pure Transformers

Pure Transformer based architectures consist of only ViT layers and have seen fewer applications in medical image segmentation compared to hybrid architectures as both global and local information is crucial for dense prediction tasks like segmentation [96]. Recently, Karimi *et al.* [153] propose a pure Transformer-based model for 3D medical image segmentation by leveraging self-attention [17] between neighboring linear embedding of 3D medical image patches. They also propose a method to effectively pre-train their

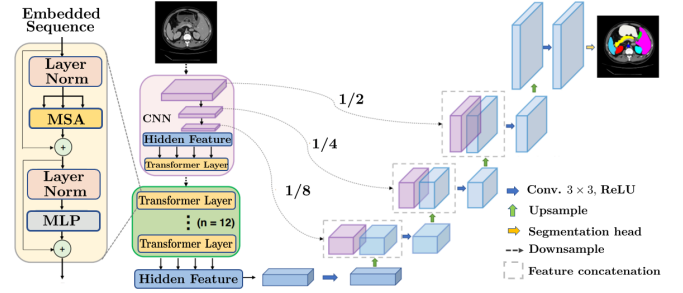


Figure 8: Overview of TransUNet architecture [96] proposed for multi-organ segmentation. It is one of the first transformer-based architecture proposed for medical image segmentation and merits both transformer and UNet. It employs a hybrid CNN-Transformer architecture for encoder, followed by multiple upsampling layers in decoder to output final segmentation mask. Image adapted from [96].

model when only a few labeled images are available. Extensive experiments show the effectiveness of their convolution-free network on three benchmark 3D medical imaging datasets related to brain cortical plate [154], pancreas, and hippocampus. One of the drawbacks of using Pure Transformer-based models in segmentation is the quadratic complexity of self-attention with respect to the input image dimensions. This can hinder the ViTs applicability in the segmentation of high-resolution medical images. To mitigate this issue, Cao *et al.* [125] propose Swin-UNet that, like Swin Transformer [126], computes self-attention within a local window and has linear computational complexity with respect to the input image. Swin-UNet also contains a patch expanding layer for upsampling decoder’s feature maps and shows superior performance in recovering fine details compared to bilinear upsampling. Experiments on Synapse and ACDC [155] dataset demonstrate the effectiveness of the Swin-UNet architectural design.

3.2.2 Hybrid Architectures

Hybrid architecture-based approaches combine the complementary strengths of Transformers and CNNs to effectively model global context and capture local features for accurate segmentation. We have further categorized these hybrid models into single and multi-scale approaches.

3.2.2.1 Single-Scale Architectures: These methods process the input image information at one scale only and have seen widespread applications in medical image segmentation due to their low computational complexity compared to multi-scale architectures. We can sub-categorized single-scale architectures based on the position of the Transformer layers in the model. These sub-categories include *Transformer in Encoder*, *Transformer between Encoder and Decoder*, *Transformer in Encoder and Decoder*, and *Transformer in Decoder*.

Transformer in Encoder. Most initially developed Transformer-based medical image segmentation approaches have Transformer layers in the model’s encoder. The first work in this category is TransUNet [96] that consists of 12 Transformer layers in the encoder as shown in Figure 8. These Transformer layers encode the tokenized image patches from the CNN layers. The resulting encoded features

are upsampled via up-sampling layers in the decoder to output the final segmentation map. With skip-connection incorporated, TransUNet sets new records (at the time of publication) on synapse multi-organ segmentation dataset [156] and automated cardiac diagnosis challenge (ACDC) [155]. In other work, Zhang *et al.* propose TransFuse [157] to effectively fuse features from the Transformer and CNN layers via BiFusion module. The BiFusion module leverages the self-attention and multi-modal fusion mechanism to selectively fuse the features. Extensive evaluation of TransFuse on multiple modalities (2D and 3D), including Polyp segmentation, skin lesion segmentation, Hip segmentation, and prostate segmentation, demonstrate its efficacy. Both TransUNet [96] and TransFuse [157] require pre-training on ImageNet dataset [158] to effectively learn the positional encoding of the images. To learn this positional bias without any pre-training, Valanarasu *et al.* [128] propose a modified gated axial attention layer [159] that works well on small medical image segmentation datasets. Furthermore, to boost segmentation performance, they propose a Local-Global training scheme to focus on the fine details of input images. Extensive experimentation on brain anatomy segmentation [160], gland segmentation [161], and MoNuSeg (microscopy) [162] demonstrate the effectiveness of their proposed gated axial attention module.

In another work, Tang *et al.* [163] introduce Swin UNETR, a novel self-supervised learning framework with proxy tasks to pre-train Transformer encoder on 5,050 images of CT dataset. They validate the effectiveness of pre-training by fine-tuning the Transformer encoder with a CNN-based decoder on the downstream task of MSD and BTCV segmentation datasets. Similarly, Sobirov *et al.* [164] show that transformer-based models can achieve comparable results to state-of-the-art CNN-based approaches on the task of head and neck tumor segmentation. Few works have also investigated the effectiveness of Transformer layers by integrating them into the encoder of UNet-based architectures in a plug-and-play manner. For instance, Cheng *et al.* [165] propose TransClaw UNet by integrating Transformer layers in the encoding part of the Claw UNet [166] to exploit multi-scale information. TransClaw-UNet achieves an absolute gain of 0.6 in dice score compared to Claw-UNet on Synapse multi-organ segmentation dataset and shows excellent generalization. Similarly, inspired from the LeViT [167], Xu *et al.* [168] propose LeViT-UNet which aims to optimize the trade-off between accuracy and efficiency. LeViT-UNet is a multi-stage architecture that demonstrates good performance and generalization ability on Synapse and ACDC benchmarks.

Transformer between Encoder and Decoder. In this category, Transformer layers are between the encoder and decoder of a U-Shape architecture. These architectures are more suitable to avoid the loss of details during down-sampling in the encoder layers. The first work in this category is TransAttUNet [169] that leverages guided attention and multi-scale skip connection to enhance the flexibility of traditional UNet. Specifically, a robust self-aware attention module has been embedded between the encoder and decoder of UNet to concurrently exploit the expressive abilities of global spatial attention and transformer self-attention. Extensive experiments on five benchmark medical imaging segmentation datasets demonstrate the




Figure 9: Overview of the interleaved encoder not-another transFormer (nnFormer) [144] for volumetric medical image segmentation. Note that convolution and transformer layers are interleaved to give full play to their strengths. Image taken from [144].

effectiveness of TransAttUNet architecture. Similarly, Yan *et al.* [170] propose Axial Fusion Transformer UNet (AFter-UNet) that contains a computationally efficient axial fusion layer between encoder and decoder to effectively fuse inter and intra-slice information for 3D medical image segmentation. Experimentation on BCV [171], Thorax-85 [172], and SegTHOR [173] datasets demonstrate the effectiveness of their proposed fusion layer.

Transformer in Encoder and Decoder. Few works integrate Transformer layers in both encoder and decoder of a U-shape architecture to better exploit the global context for medical image segmentation. The first work in this category is UTNet that efficiently reduces the complexity of the self-attention mechanism from quadratic to linear [174]. Furthermore, to model the image content effectively, UTNet exploits the two-dimensional relative position encoding [20]. Experiments show strong generalization ability of UTNet on multi-label and multi-vendor cardiac MRI challenge dataset cohort [175]. Similarly, to optimally combine convolution and transformer layers for medical image segmentation, Zhou *et al.* [144] propose nnFormer, an interleave encoder-decoder based architecture, where convolution layer encodes precise spatial information and Transformer layer encodes global context as shown in Fig. 9. Like Swin Transformers [126], the self-attention in nnFormer has been computed within a local window to reduce the computational complexity. Moreover, deep supervision in the decoder layers has been employed to enhance performance. Experiments on ACDC and Synapse datasets show that nnFormer surpass Swin-UNet [125] (transformer-based medical segmentation approach) by over 7% (dice score) on Synapse dataset. In other work, Lin *et al.* propose Dual Swin Transformer UNet (DS-TransUNet) [176] to incorporate the advantages of Swin Transformer in U-shaped architecture for medical image segmentation. They split the input image into non-overlapping patches at two scales and feed them into the two Swin Transformer-based branches of the encoder. A novel Transformer Interactive Fusion module has been proposed to build long-range dependencies between different scale features in encoder. DS-TransUNet outperforms CNN-based methods on four standard datasets related to Polyp segmentation, ISIC 2018, GLAS, and Datascience bowl 2018.

Transformer in Decoder. Li *et al.* [177] investigate the use of Transformer as an upsampling block in the decoder of the UNet for medical image segmentation. Specifically, they adopt a window-based self-attention mechanism to better complement the upsampled feature maps while maintaining

Method	Organ	Modality	Type	Datasets	Metrics	Arch.	P.T.	Highlights
TransUNet [96]	Multi-organ	CT, MRI	2D	Synapse [156], ACDC [155]	Dice, Hausdorff distance	Hybrid	Yes	Encodes strong global context by treating the image features as sequences but also well utilizes the low-level CNN features via a u-shaped hybrid architectural design.
TransFuse [157]	Multi-organ	Colonoscopy	2D, 3D	KVASIR [178], Clinic DB [179], Colon DB [180], EndoScene [181], ETIS [182], ISIC 2017 [103], MSD [177]	Dice, Jaccard index	Hybrid	Yes	Leverages the inductive bias of CNNs on modeling spatial correlation and the powerful capability of Transformers on modelling global relationship. Novel Bi-Fusion module to fuse CNN and Transformers features for segmentation.
MedT [183]	Multi-organ	Ultrasound, Microscopy	2D	Brain US (Private), GLAS [161], MoNuSeg [184]	F1	Hybrid	No	Gated axial attention layer for positional encoding. Local global training for training on both full resolution images as well in patches.
Conv. Free [153]	Multi-organ	MRI, CT	3D	Brain cortical, (private)	Dice, Hausdorff distance, Average Symmetric Surface Distance	Pure	Yes	Convolutional free medical segmentation model, based on self-attention between neighboring 3D patches.
CoTr [112]	Multi-organ	CT	3D	BCV [171]	Dice	Hybrid	Yes	Deformable self-attention mechanism to reduce the computational and spatial complexities of modelling the long range dependency.
SpecTr [116]	Bile-duct	Hyperspectral	3D	CholeDock [47]	Dice, IoU, Hausdorff distance	Hybrid	No	First application to hyperspectral and learn contextual feature across spectral dimension.
TransBTS [138]	Brain	MRI	3D	BraTS 19 [185], BraTS 20 [186]	Dice, Hausdorff distance	Hybrid	Yes	3D CNN for capturing local volumetric features and transformers for encoding global features.
U-Transformer [187]	Multi-organ	CT	2D	TCIA [188], Private multi-organ	Dice	Hybrid	No	Propose self and cross-attention modules to model long-range interactions and spatial dependencies.
UNETR [35]	Brain, Spleen	MRI, CT	3D	BITC, MSD [171]	Dice, Hausdorff distance	Hybrid	No	Transformer as the encoder to learn sequence representations of the input volume and effectively capture the global multi-scale information
PMTrans [183]	Multi-organ	Microscopy, CT, PET	2D	GLAS [161], MoNuSeg [184], HECKTOR [189]	Dice	Hybrid	No	Integrate multi-scale attention and CNN feature extraction using a pyramidal network architecture. An adaptive partitioning scheme was implemented to retain informative relations and to access different receptive fields efficiently.
Swin-UNet [125]	Multi-organ	CT	2D	Synapse [156], ACDC [155]	Dice	Pure	Yes	Swin transformer based pure transformer architecture to segmentation with patch expanding layer design in decoder.
Segtran [190]	Multi-organ	Fundus, Colonoscopy, MRI	2D, 3D	REFUGE 20 [191], BraTS 19 [185], CVC [192], KVASIR [178]	Dice	Hybrid	Yes	Squeeze-and-Expansion transformer where squeeze block regularize the self-attention module and expansion block learns diversified representations.
MBTNet [136]	Eye	Pathology	2D	TMEM3000 (private), Alisairne [137]	Dice, F1, Sensitivity, Specificity	Hybrid	No	Body edge branch for precise edge location information for corneal endothelial cells segmentation.
DS-TransUNet [176]	Multi-organ	Colonoscopy, Histology	2D	Kvasir [178], Colon DB [180] and Clinic DB [179], EndoScene [181], ETIS [182], ISIC 18 [101], GLAS [161], Data Science Bowl 18 [193]	Mean Dice, Mean IoU, Precision, Recall	Hybrid	Yes	Dual-branch Swin Transformer in encoder and decoder to extract multiscale representation. Transformer Interactive Fusion module to build long-range dependencies between features of different scales
MCTrans [39]	Multi-organ	Colonoscopy, Pathology	2D	Pannuke dataset, Colon DB [180], Clinic DB [179], ETIS [182], KVASIR [178], ISIC 2018 [101]	Dice	Hybrid	No	Transformer self-attention for cross-scale contextual dependencies Transformer cross attention layer for semantic correspondence.
Li et al. [177]	Multi-organ	MRI, CT	2D	Synapse [156], MSD Brain [171]	Dice, Hausdorff distance	Hybrid	No	Investigate the use of transformer decoder for medical image segmentation and its usage in upsampling.
UNet [34]	Heart	MRI	2D	MRI Challenge Cohort [175]	Dice, Hausdorff distance	Hybrid	No	Self-attention modules in encoder and decoder. Design relative position encoding to reduce the complexity of self-attention from quadratic to linear.
TransClaw UNet [165]	Multi-organ	CT	2D	Synapse [156]	Dice, Hausdorff distance	Hybrid	No	Integrated transformer layer in the encoder path of Claw-UNet to extract shallow spatial features.
TransAttUNet [169]	Multi-organ	Xray, CT	2D	ISIC 2018 [101], JSRT [194], Montgomery [195], NIH [196], Clean-CC-CII [197], Data Science Bowl 18 [193], GLAS [161]	Dice, F1	Hybrid	No	Multi-level guided attention and multi-scale skip connections to mitigate information recession problem.
LeViT-UNet [168]	Multi-organ	CT, MRI	2D	Synapse [156], ACDC [155]	Dice, Hausdorff distance	Hybrid	Yes	Integrate multiscale LeViT architecture as the a encoder in UNet.
Polyp-PVT [198]	Multi-organ	Colonoscopy	2D	KVASIR [178], Clinic DB [179], Colon DB [180], EndoScene [181], ETIS [182], KIT521 Challenge [119]	Dice, IoU, MAE, Weighted F-measure, S-measure, E-measure, Dice, Surface Dice	Hybrid	No	Pyramid vision transformer backbone as encoder to extract robust features. Proposed architectural components to handle noise, occlusions, and capturing global semantic cues.
CoTRNet [117]	Kidney	CT	2D		Dice, Surface Dice	Hybrid	Yes	CNN and transformer based interleaved encoder-decoder. Supervision of decoder's hidden layers.
nnFormer [144]	Multi-organ	CT, MRI	3D	Synapse [156], ACDC [155]	Dice	Hybrid	Yes	Interleaved convolution and self-attention based encoder-decoder architecture.
MISSFormer [199]	Multi-organ	CT, MRI	2D	Synapse [156], ACDC [155]	Dice, Hausdorff distance	Hybrid	No	Hierarchical encoder-decoder network with enhanced transformer block to mitigate the problem of feature inconsistency
TransBridge [109]	Heart	Echocardiography	2D	EchoNet-Dynamic [111]	Dice, Hausdorff distance	Hybrid	No	Shuffling layer and group convolution for patch embedding to significantly reduce the number of parameters.
BiTr-UNet [139]	Brain	MRI	3D	BraTS 21 [140]	Dice, Hausdorff distance	Hybrid	No	Refined version of TransBTS with two sets of ViT layers instead of one.
GT UNet [105]	Tooth	X-ray	2D	Tooth root dataset (private)	Dice, Accuracy, Sensitivity, Specificity, Jaccard similarity	Hybrid	No	Group transformer layers to reduce computational cost. Fourier descriptor based loss function to integrate shape prior.
BAT [97]	-	Dermoscopy	2D	ISIC 2016+PH2 [99], ISIC 2018 [101]	Dice, IoU	Hybrid	Yes	Boundary-wise attention gate is added at the end of each transformer encoder layer to tackle challenging cases with ambiguous boundaries.
AFTER-UNet [170]	Multi-organ	CT	3D	BCV [171], Thorax-85 [172], SegTHOR [173]	Dice	Hybrid	No	Axial fusion mechanism to fuse intra-slice and inter-slice contextual information to guide segmentation.
ViT-UNet [143]	Multi-organ	CT, MRI	3D	BraTS 21 [140], MSD [171]	Dice, Hausdorff distance	Hybrid	Yes	U-shaped encoder-decoder design. Encoder has two consecutive self-attention layers to encode local and global cues, and our decoder has novel parallel shifted window based self and cross attention blocks to capture fine details.
Swin UNETR [145]	Brain	MRI	3D	BraTS 21 [140]	Dice, Hausdorff distance	Hybrid	Yes	Swin UNet based architecture that consists of Swin transformer as the encoder and a CNN-based decoder. Computes self-attention in an efficient shifted window partitioning scheme.

Table 3: An overview of ViT-based approaches for medical image segmentation. P.T: pretraining.



Figure 10: Qualitative results of brain tumor segmentation task using transformer. From left to right: Ground truth image, UNETR [35] (ViT-based), TransBTS [138] (ViT-based), CoTr [200] (ViT-based), and UNet [114] (CNN based). Note that transformer-based approaches demonstrate better performance in capturing the fine-grained details of brain tumors as compared to CNN-based method. Image courtesy [35].

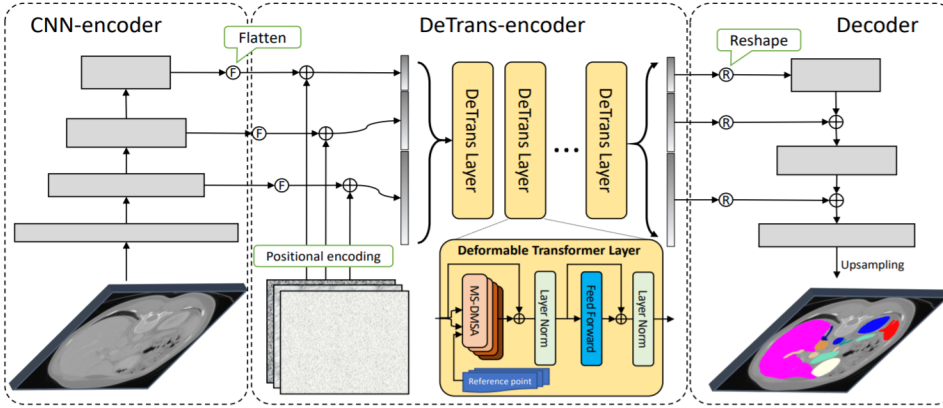


Figure 11: Overview of CNN and a Transformer (CoTr) architecture [112] proposed for 3D medical image segmentation. It consists of CNN encoder (left) to extract multi-scale features from the input, followed by DeTrans-encoder (yellow blocks) to process the flattened multi-scale feature maps. Output features from encoder are fed to the CNN decoder (right) for segmentation mask prediction. Image courtesy [112].

efficiency. Experiments on MSD Brain and Synapse datasets demonstrate the superiority of their architecture compared to bilinear upsampling. In another work, Li *et. al* [190] propose SegTran, a Squeeze-and-Expansion Transformer for 2D and 3D medical image segmentation. Specifically, the squeeze block regularizes the attention matrix, and the expansion block learns diversified representations. Furthermore, a learnable sinusoidal positional encoding has been proposed that helps the model to encode spatial relationships. Extensive experiments on Polyp, BraTS19, and REFUGE20 (fundus images) segmentation challenges demonstrate the strong generalization ability of Segtran.

3.2.2.2 Multi-Scale Architectures: These architectures process input at multiple scales to effectively segment organs having irregular shapes and different sizes. Here, we highlight various attempts to integrate the multi-scale architectures for medical image segmentation. We further group these approaches into 2D and 3D segmentation categories based on the input image type.

2D Segmentation. Most ViT-based multi-organ segmentation approaches struggle to capture information at multiple scales as they partition the input image into fixed-size patches, thereby losing useful information. To address this issue, Zhang *et. al.* [183] propose a pyramid medical transformer, PMTrans, that leverage multi-resolution attention to capture correlation at different image scales using a pyramidal architecture [201]. PMTrans works on multi-resolution images via an adaptive partitioning scheme of patches to access different receptive fields without changing the overall complexity of self-attention computation. Ex-

tensive experiments on three medical imaging datasets of GLAS [161], MoNuSeg [184], and HECKTOR [189] show the effectiveness of exploiting multi-scale information. In other work, Ji *et. al.* [39] propose a Multi-Compound transformer (MCTrans) that learns not only feature consistency of the same semantic categories but also capture correlation among different semantic categories for accurate segmentation [202]. Specifically, MCTrans captures cross-scale contextual dependencies via the Transformer self-attention module and learned semantic correspondence among different categories via Transformer Cross-Attention module. An auxiliary loss has also been introduced to improve feature correlation of the same semantic category. Extensive experiments on six benchmark segmentation datasets demonstrate the effectiveness of the architectural components of MCTrans.

3D Segmentation. The majority of multi-scale architectures have been proposed for 2D medical image segmentation. To directly handle volumetric data, Hatamizadeh *et. al.* [35] propose a ViT-based architecture (UNETR) for 3D medical image segmentation. UNETR consists of a pure transformer as the encoder to learn sequence representations of the input volume. The encoder is connected to a CNN-based decoder via skip connections to compute the final segmentation output. UNETR achieves impressive performance on BTCV [203] and MSD [171] segmentation datasets as shown in Fig. 10. One of the drawbacks of UNETR is its large computational complexity in processing large 3D input volumes. To mitigate this issue, Xie *et. al* [112] propose a computationally efficient deformable self-attention module [204] that casts attention only to a small set using




Figure 12: Dice results of BTCV challenge on multi-organ segmentation for various transformer based methods. It can be seen that Swin-UNETR is able to achieve on average 13% improvement in dice coefficient score compare to SETR method, indicating rapid pace of research in the field. Transformer based approaches used for the comparison include SETR NUP [25], SETR PUP [25], SETR MLA [25], TransUNet [96], CoTr* [112] (small CNN encoder compared to CoTr), CoTr [112], UNETR [35], and Swin UNETR [163]. Note: Avg: Average results (over 12 organs), AG: left and right adrenal glands, Pan: pancreas, Sto: stomach, Spl: spleen, Liv: liver, Gall: gallbladder.

multi-scale features, as shown in Figure 11, to reduce the computational and spatial complexities. Experiments on BTCV [203] demonstrate the effectiveness of their deformable self-attention module for 3D multi-organ segmentation.

3.3 Discussion

From the extensive literature reviewed in this section, we note that the medical image segmentation area is heavily impacted by transformer-based models, with more than 50 publications within one year since the inception of the first ViT model [22]. We believe such interest is due to the availability of large medical segmentation datasets and challenge competitions associated with them in top conferences compared to other medical imaging applications. As shown in Fig. 12, a recent transformer-based hybrid architecture is able to achieve 13% performance gain in terms of dice score compared to the baseline transformer model, indicating rapid progression of the field. In short, ViT-based architectures have achieved impressive results over benchmark medical datasets, competing and most of the time improving over CNN-based segmentation approaches (see Table 3 for details). Below, we briefly describe some of the challenges associated with ViTs based medical segmentation methods and give possible solutions based on insights from the relevant papers discussed.

As mentioned before, the high computational cost associated with extracting features at multiple levels hinders the applicability of multi-scale architectures in medical segmentation tasks. These multi-scale architectures exploit processing input image information at multiple levels and achieve superior performance than single-scale architectures. Therefore, designing efficient transformer architectures for multi-scale processing requires more attention.

Most of the proposed ViT-based models are pre-trained on the ImageNet dataset for the downstream task of medical image

segmentation. This approach is sub-optimal due to the large domain gap between natural and medical image modalities. Recently, few attempts have been made to investigate the impact of self-supervised pre-training on medical imaging datasets on the ViTs segmentation performance. However, these works have shown that ViT pre-trained on one modality (CT) gives unsatisfactory performance when applied directly to other medical imaging modalities (MRI) due to the large domain gap making it an exciting avenue to explore. We defer detailed discussion related to pre-training ViTs for downstream medical imaging tasks to Sec. 11.1.

Moreover, recent ViT-based approaches mainly focus on 2D medical image segmentation. Designing customized architectural components by incorporating temporal information for efficient high-resolution and high-dimensional segmentation of volumetric images has not been extensively explored. Recently, few efforts have been made, e.g., UNETR [35] uses Swin Transformer [126] based architectures to avoid quadratic computing complexity; however, it requires further attention from the community.

In addition to focusing on the scale of datasets, with the advent of ViTs, we note there is a need to collect more diverse and challenging medical imaging datasets. Although diverse and challenging datasets are also crucial to gauge the performance of ViTs in other medical imaging applications, they are particularly relevant for medical image segmentation due to a major influx of ViT-based models in this area. We believe these datasets will play a decisive role in exploring the limits of ViTs for medical image segmentation.

4 MEDICAL IMAGE CLASSIFICATION

Accurate classification of medical images plays an essential role in aiding clinical care and treatment. In this section, we comprehensively cover applications of ViTs in medical image classification. We have broadly categorized these approaches into COVID-19, tumor, and retinal disease classification based methods due to a different set of challenges associated with these categories as shown in Fig. 13.

4.1 COVID-19 Diagnosis

Studies suggest that COVID-19 can potentially be better diagnosed with radiological imaging as compared to tedious real-time polymerase chain reaction (RT-PCR) test [205]–[207]. Recently, ViTs have been successfully employed for diagnosis and severity prediction of COVID-19, showing SOTA performance. In this section, we briefly describe the impact of ViTs in advancing recent efforts on automated image analysis for the COVID-19 diagnosis process. Most of these works use three modalities, including Computerized tomography (CT), Ultrasound scans (US), and X-ray. We have further categorized ViT-based COVID-19 classification approaches into *Black-box models* and *Interpretable models* according to the level of explainability offered.

4.1.1 Black-Box Models

ViT-based Black-box models for COVID-19 imaging classification generally focus on improving accuracy by designing novel and efficient ViT architectures. However, these models are not easily interpretable, making it challenging to gain user-trust. We have further sub-categorized black-box models into 2D and 3D categories, depending on the input image type. Below, we briefly describe these approaches:

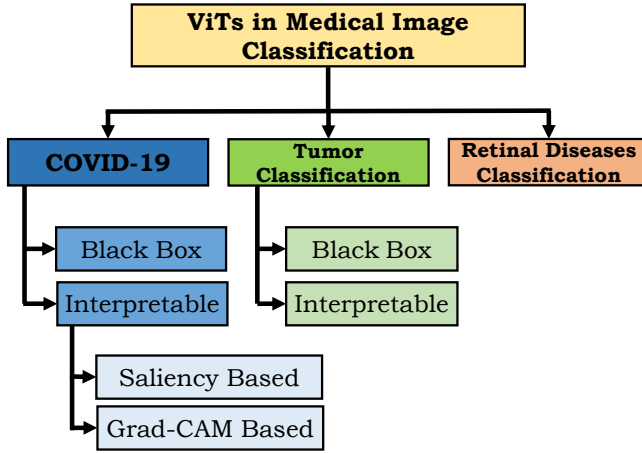


Figure 13: Taxonomy of ViT-based medical image classification approaches. The influx of ViT-based COVID-19 classification approaches makes it a dominating category in the taxonomy.

2D: The High computational cost of ViTs hinders their deployment on portable devices, thereby limiting their applicability in real-time COVID-19 diagnosis. Perera *et al.* [208] propose a lightweight **Point-of-Care Transformer (POCFormer)** to diagnose COVID-19 from lungs images captured via portable devices. Specifically, POCFormer leverages Linformer [174] to reduce the space and time complexity of self-attention from quadratic to linear. POCFormer has two million parameters that are about half of MobileNetv2 [209], thus making it suitable for real-time diagnosis. Experiments on COVID-19 lungs POCUS dataset [210], [211] demonstrate the effectiveness of their proposed architecture with above 90% classification accuracy. In other work, Liu *et al.* [212] proposed ViT-based model for COVID-19 diagnosis by exploiting a new attention mechanism named Vision Outlooker (VOLO) [213]. VOLO is effective for **encoding fine-level features** into ViT token representation, thereby improving classification performance. Further, they leverage the transfer learning approach to handle the issue of insufficient and generally unbalanced COVID-19 datasets. Experiments on two publicly available COVID-19 CXR datasets [211], [214] demonstrate the effectiveness of their architecture. Similarly, Jiang *et al.* [215] leverage **Swin Transformer** [126] and **Transformer-in-Transformer** [216] to classify COVID-19 images from Pneumonia and normal images. To further boost the accuracy, they employ model ensembling using a weighted average. Research progress in ViT-based COVID-19 diagnosis approaches is heavily impeded due to the requirement of a large amount of labeled COVID-19 data, thereby demanding collaborations among hospitals. This collaboration is difficult due to limited consent by patients, privacy concerns, and ethical data usage [217]. To mitigate this issue, Park *et al.* [218] proposed a **Federated Split Task-Agnostic (FESTA) framework that leveraged the merits of Federated and Split Learning** [219], [220] in utilizing ViT to simultaneously process multiple chest X-ray tasks, including the diagnosis in COVID-19 Chest X-ray images on a massive decentralized dataset. Specifically, they split ViT into the shared transformer body and task-

specific heads and demonstrate the suitability of ViT body to be shared across relevant tasks by leveraging multitask-learning (MTL) [221] strategy as shown in Fig. 14. They affirm the suitability of ViTs for collaborative learning in medical imaging applications via extensive experiments on the CXR dataset.

3D: Most of the ViT-based approaches for COVID-19 classification operate on 2D information only. However, as suggested by Kwee *et al.* [222], the symptoms of COVID-19 might be present at different depths (slices) for different patients. To exploit both 2D and 3D information, Hsu *et al.* [223] propose a hybrid network consisting of transformers and CNNs. Specifically, they determine the importance of slices based on significant symptoms in the CT scan via Wilcoxon signed-rank test [224] with Swin Transformer [126] as backbone network. To further exploit the intrinsic features in the spatial and temporal dimensions, they propose a Convolutional CT Scan Aware Transformer module to fully capture the context of the 3D scans. Extensive experiments on the COVID-19-CT dataset show the effectiveness of their proposed architectural components. Similarly, Zhang *et al.* [225], [226] also proposed Swin Transformer based two-stage framework for the diagnosis of COVID-19 in the 3D CT scan dataset [227]. Specifically, their framework consists of UNet based lung segmentation model followed by the image classification with Swin Transformer [126] backbone.

4.1.2 Interpretable Models

Interpretable models aim to show the features that influence the decision of a model the most, generally via visualization techniques like saliency-based methods, Grad-CAM, etc. Due to their interpretable nature, these models are well suited to gain the trust of physicians and patients and therefore have paved their way for clinical deployment. We have further divided interpretable models into saliency-based [228] and Grad-CAM [229] based visualization approaches.

Saliency Based Visualization. Park *et al.* [231], propose a ViT-based method for COVID-19 diagnosis by exploiting the low-level CXR features extracted from the pre-trained backbone network. The backbone network has been trained in a self-supervised manner (using contrastive-learning based SimCLR [232] method) to extract abnormal CXR features embeddings from large and well-curated CXR dataset of CheXpert [233]. These feature embeddings have been leveraged by ViT model for high-level diagnosis of COVID-19 images. Extensive experiments on three CXR test datasets acquired from different hospitals demonstrate the superiority of their approach compared to CNN-based models. They also validated the generalization ability of their proposed approach and adopted saliency map visualizations [234] to provide interpretable results. Similarly, Gao *et al.* [235] propose COVID-ViT to classify COVID from non-COVID images as part of the MIA-COVID19 challenge [227]. Their experiments on 3D CT lungs images demonstrated the superiority of ViT-based approach over DenseNet [236] baseline in terms of F1 score. In another work, Mondal *et al.* [230] introduce xViTCOS for COVID-19 screening from lungs CT and X-ray images. Specifically, they pre-train xViTCOS on ImageNet to learn generic image representations and fine-tune the pre-trained model on a large chest radiographic dataset. Further, xViTCOS leverage the explainability-driven

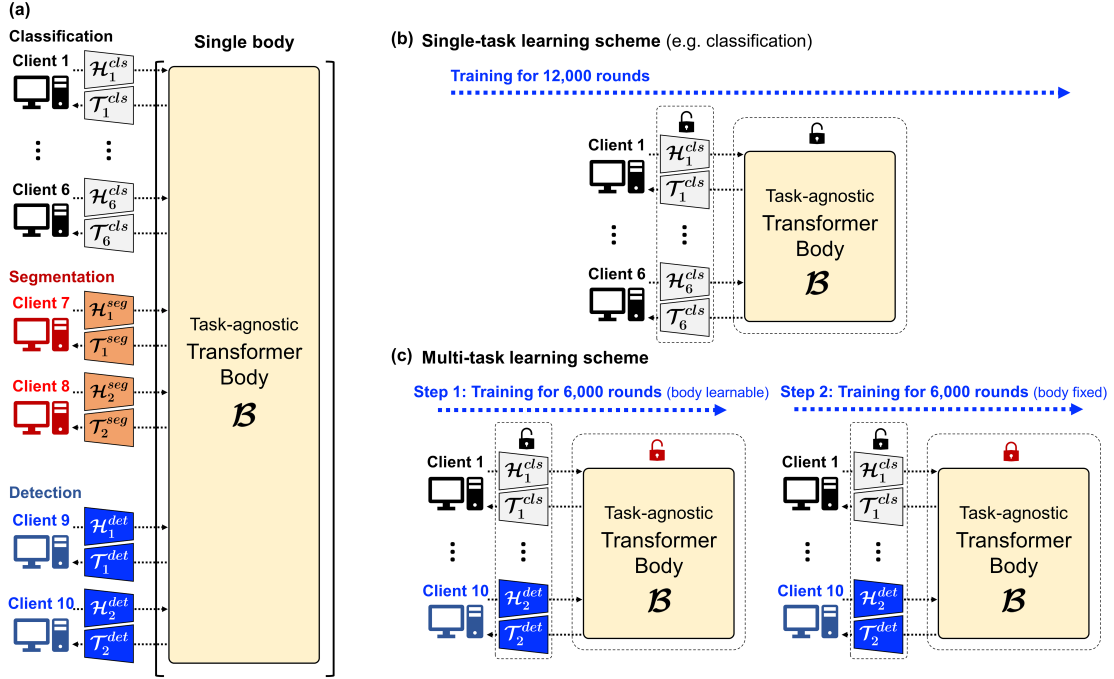


Figure 14: Implementation details of Federated Split Task-Agnostic (FESTA) framework [218] equipped with a Transformer to simultaneously process multiple chest X-ray tasks including diagnosis of COVID-19. (a) Experimental setting for multi-task learning of classification, segmentation, and detection of chest X-ray images. The clients only train the head (\mathcal{H}) and tail (\mathcal{T}) parts of the network, whereas the transformer body (\mathcal{B}) is shared across multiple clients. In the second step, the embedded features in the head are utilized by transformers for processing of individual tasks. (b) shows training scheme for single task. (c) shows training scheme for multi-task learning. Image courtesy [218].

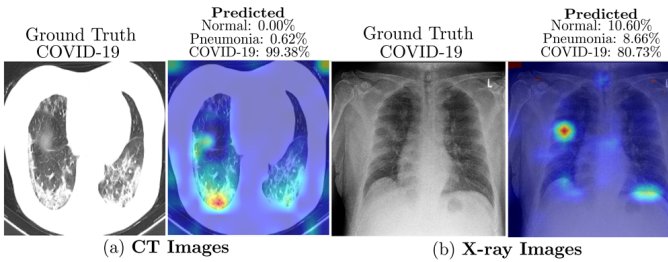


Figure 15: CT scans (a) and X-ray (b) images along with their ground truth labels (left) and saliency maps (right). For figure (a), xViTCOS-CT localized suspicious lesion regions exhibiting ground glass opacities, consolidation, reticulations in bilateral postero basal lung. xViTCOS-CT [230] is able predict these regions correctly. For figure (b), radiologist's interpretation is: *thick walled cavity in right middle zone with surrounding consolidation*. As shown in last column, xViTCOS-CXR [230] is able predict it correctly. Figure courtesy [230].

saliency-based approach [234] with clinically interpretable visualizations to highlight the role of critical factors in the resulting predictions, as shown in Figure 15. Experiments on COVID CT-2A [237] and their privately collected Chest X-ray dataset demonstrate the effectiveness of xViTCOS.

Grad-CAM Based Visualization. Shome *et al.* [238] propose a ViT-based model to diagnose COVID-19 infection at scale. They combine several open-source COVID-19 CXR datasets to form a large-scale multi-class and binary classification dataset. For better visual representation and

model interpretability, they further create Grad-CAM based visualization [229].

4.2 Tumor Classification

A tumor is an abnormal growth of body tissues and can be cancerous (malignant) or noncancerous (benign). Early-stage malignant tumor diagnosis is crucial for subsequent treatment planning and can greatly improve the patient's survival rate. In this section, we review ViT-based models for tumor classification. These models can be mainly categorized into *Black-box models* and *Interpretable models*. We highlight the relevant anatomies in bold.

Black-Box Models. TransMed [240] is the first work that leverages ViTs for medical image classification. It is a hybrid CNN and transformer-based architecture that is capable of classifying **parotid** tumors in the multi-modal MRI medical images. TransMed also employs a novel image fusion strategy to effectively capture mutual information from images of different modalities, thereby achieving competitive results on their privately collected parotid tumor classification dataset. Later, Lu *et al.* [241] propose a two-stage framework that first performs contrastive pre-training on glioma sub-type classification in the **brain** followed by the feature aggregation via proposed transformer-based sparse attention module. Ablation studies on TCGA-NSCLC [242] dataset show the effectiveness of their two-stage framework. For the task of **breast** cancer classification, Gheflati *et al.* [243] systematically evaluate the performance of pure and hybrid pre-trained ViT models. Experiments on two breast ultrasound datasets provided by Al-Dhabyani *et al.* [244]




Figure 16: Overview of Transformer based Multiple Instance Learning (TransMIL) architecture [239] for whole slide brain tumor classification. Patches of WSI are embedded in the feature space of ResNet-50. The sequence of embedded features are then processed by their proposed pipeline that include: squaring of sequence, Correlation modelling of the sequence, conditional position encoding (via Pyramid Position Encoding Generator (PPEG) module) and local information fusion, feature aggregation, and mapping from transformer space to label space. Image taken from [239].




Figure 17: Left: The area within the blue region is the cancer region. Middle: Attention scores from TransMIL are visualised as a heatmap (red for tumor and blue for normal) to interpret the important morphology used for diagnosis. Right: Zoomed-in view of the black square in the middle figure. Figure courtesy [239].

and Yap *et al.* [245] shows that ViT-based models provide better results than those of the CNNs for classifying images into benign, malignant, and normal categories. Similarly, other works employ hybrid Transformer-CNN architectures to solve medical classification problem for different organs. For instance, Khan *et al.* [246] propose Gene-Transformer to predict the **lung** cancer subtypes. Experiments on TCGA-NSCLC [242] dataset demonstrates the superiority of Gene Transformer over CNN baselines. Chen *et al.* [247] present a multi-scale GasHis-Transformer to diagnose gastric cancer in the **stomach**. Jiang *et al.* [248] propose a hybrid model to diagnose acute lymphocytic leukemia by using symmetric cross-entropy loss function.

Interpretable Models. Since the annotation procedure is expensive and laborious, one label is assigned to a set of instances (bag) in whole slide imaging (WSI) based pathology diagnosis. This type of weakly supervised learning is known as Multiple Instance Learning [249], where a bag is labeled positive if at least one instance is positive or labeled negative when all instances in a bag are negative. Most of the current MIL methods assume that the instances in each bag are independent and identically distributed, thereby neglecting the correlation among different instances. Shao *et al.* [239] present TransMIL to explore both morphological and spatial information in weakly supervised WSI classification. Specifically, TransMIL aggregates morphological information

with two transformer-based modules and a position encoding layer as shown in Fig. 16. To encode spatial information, a pyramid position encoding generator is proposed. Further, the attention scores from the TransMIL have been visualized to demonstrate interpretability, as shown in Fig. 17. TransMIL shows state-of-the-art performance on three different computational pathology datasets CAMELYON16 (**breast**) [250], TCGA-NSCLC (**lung**) [242], and TCGA-R (**kidney**) [251]. To diagnose **lung** tumors, Zheng *et al.* [252] propose graph transformer network (GTN) to leverage the graph-based representation of WSI. GTN consists of a graph convolutional layer [253], a transformer layer, and a pooling layer. GTN further employs GraphCAM [234] to identify regions that are highly associated with the class label. Extensive evaluations on TCGA dataset [242] show the effectiveness of GTN.

4.3 Retinal Disease Classification

Yu *et al.* [258] propose MIL-ViT model which is first pre-trained on a large fundus image dataset and later fine-tuned on the downstream task of the retinal disease classification. MIL-ViT architecture uses MIL-based head that can be used with ViT in a plug-and-play manner. Evaluation performed on APTOS2019 [255] and RFMiD2020 [259] datasets shows that MIL-ViT is achieving more favorable performance than CNN-based baselines. Most data-driven approaches treat diabetic retinopathy (DR) grading and lesion discovery as two separate tasks, which may be sub-optimal as the error may propagate from one stage to the other. To jointly handle both these tasks, Sun *et al.* [260] propose lesion aware transformer (LAT) that consists of a pixel relation based encoder and a lesion-aware transformer decoder. In particular, they leverage transformer decoder to formulate lesion discovery as a weakly supervised lesion localization problem. LAT model sets state-of-the-art on Messidor-1 [261], Messidor-2 [261], and EyePACS [262] datasets. Yang *et al.* [273] propose a hybrid architecture consisting of convolutional and Transformer layers for fundus disease classification on OIA dataset [274]. Similarly, Wu *et al.* [275] and Aldahou *et al.* [276] also verify that ViT models are more accurate in DR grading than their CNNs counterparts.

Method	Organ	Modality	Type	Datasets	Metrics	Arch.	Highlights
TransMed [240]	Ear	MRI (T1,T2)	3D	MRI (private)	Accuracy Precision	Pure	First ViT-based multi-modal medical image classification approach with novel multi-modal fusion strategy
TransMIL [239]	Multi-organ	Pathology	2D	Camelyon16 [250] TCGA-NSCLC [254] TCGA-RCC [251]	Accuracy AuC	Hybrid	Transformer based architecture to explore morphological and spatial information for Whole Slide Image classification.
Matsokus <i>et al.</i> [33]	Multi-organ	Mammograms Dermoscopy	2D	APTOS-2019 [255] ISIC-2019 [256] CBIS-DDSM [257]	Recall AuC	Pure	Systemetic study of whether one should replace CNNs with ViTs for medical image classification.
Gheflati <i>et al.</i> [243]	Breast	Ultrasound	2D	BUSY [244] Yap <i>et al.</i> [245]	Accuracy AuC	Pure	First application of ViTs to ultrasound images classification.
GTN [252]	Lung	Microscopy	2D	TCGA dataset [242]	Accuracy Precision Sensitivity Specificity Recall	Hybrid	Consists of a graph convolutional layer, a transformer module, and a pooling layer for accurate classification of WSI images.
MIL-ViT [258]	Eye	Fundus	2D	APTOS-2019 [255] RFMid2020 [259]	Accuracy AuC, F1 Precision Recall	Pure	First pretrained on a large fundus image dataset and later fine-tuned on the downstream task of the retinal disease classification.
LAT [260]	Eye	Fundus	2D	Messidor-1 [261] Messidor-2 [261] EyePACS [262]	AuC Kappa	Hybrid	Formulate lesion discovery as a weakly supervised lesion localization problem via a transformer decoder. Jointly solve diabetic retinopathy grading and lesion discovery.
COVID-19							
Park <i>et al.</i> [231]	Chest	X-ray	2D	CheXpert [233]	AuC	Hybrid	Leveraging a backbone network trained to find low-level abnormal CXR findings in pre-built large-scale dataset to embed feature corpus suitable for high-level disease classification.
POCFORMER [208]	Chest	Ultrasound	2D	POCUS [210]	Recall, F1 Specificity Sensitivity Accuracy	Pure	Proposed light weight transformer architecture that demonstrates the efficiency and performance improvements.
FESTA [218]	Chest	X-ray	2D	CheXpert [233] SIM-ACR [263] RSNA [264]	Recall, F1 Specificity Sensitivity AuC	Pure	Utilize ViT to simultaneously process multiple chest X-ray tasks, including the diagnosis in COVID-19 Chest X-ray images on a massive decentralized dataset.
Liu <i>et al.</i> [212]	Chest	X-ray	2D	COVID-19-1 [214] COVID-19-2 [211]	Accuracy	Pure	Explore Volo tailored with transfer learning technique to effectively encodes fine-level features into the token representations.
COVID-ViT [235]	Chest	CT	3D	MIA-COV19 [227]	Accuracy, F1	Pure	Propose ViT based architecture to classify COVID-19 CT images in MIA-COV19 competition.
xViTCOS [230]	Chest	X-ray, CT	2D	xViTCOS-CT [237] xViTCOS-CXR [265]	Precision Recall, F1 Specificity, NPV	Pure	Propose ViT based multi-stage transfer learning technique to address the issue of COVID-19 data scarcity. The approach is clinically interpretable.
Hsu <i>et al.</i> [223]	Chest	CT	3D	COV19 CT DB [227]	Accuracy Recall, F1 Precision	Hybrid	Importance of slices are determined in CT scan via Wilcoxon signed-rank test. Then, spatial and temporal features are exploited via proposed convolutional CT scan Aware Transformer module.
Zhang <i>et al.</i> [225]	Chest	CT	3D	MIA-COV19 [227]	F1 score	Hybrid	Swin Transformer based two stage framework for diagnosis of COVID-19 in 3D CT scans.
COViT-GAN [266]	Chest	CT	2D	COVID-CT [227] Sars-CoV-2 [267]	Accuracy Sensitivity Precision, F1	Hybrid	Generate synthetic images using a self-attention generative adversarial network and use it as a data augmentation method to alleviate the problem of limited data and improve performance.
COVID-Trans. [238]	Chest	X-ray	2D	El-Shafai <i>et al.</i> [268] Sait <i>et al.</i> [269] Qi <i>et al.</i> [270]	Accuracy Precision AuC, Recall, F1	Pure	Propose a ViT model to diagnose COVID-19 at scale. Combines several open source COVID-19 chest X-ray datasets to form large dataset for binary and multi-class classification.

Table 4: An overview of ViT-based approaches for medical image classification. **Teal color** indicates that the model is interpretable.

Initialization	Model	APTOS2019, $\kappa \uparrow$	ISIC2019, Recall \uparrow
Random	ResNet50	0.849 ± 0.022	0.662 ± 0.018
	DeiT-S	0.687 ± 0.017	0.579 ± 0.028
ImageNet (supervised)	ResNet50	0.893 ± 0.004	0.810 ± 0.008
	DeiT-S	0.896 ± 0.005	0.844 ± 0.021
ImageNet (supervised) + Self-supervised with DINO [271]	ResNet50	0.894 ± 0.008	0.833 ± 0.007
	DeiT-S	0.896 ± 0.010	0.853 ± 0.009

Table 5: Comparison of vanilla CNNs vs. ViTs with different initialization strategies on *medical imaging* classification tasks. For APTOS 2019 [255] and ISIC 2019 [272] datasets quadratic Cohen Kappa and recall score have been reported. *First row*: For randomly initialized networks, CNNs outperform ViTs. *Second row*: ViTs appear to benefit significantly from pre-training on ImageNet dataset. *Third row*: Both ViTs and CNNs perform better with self-supervised pretraining. Table taken from [33].

4.4 Discussion

In this section, we provide a comprehensive overview of about 25 papers related to applications of ViTs in medical image classification.

In particular, we see a surge of Transformer-based architectures for diagnosing COVID-19, compelling us to develop taxonomy accordingly. Below, we briefly highlight some of the challenges associated with this area, identify recent trends, and provide future directions worthy of further exploration.

The lack of large COVID-19 datasets hindered the applicability of ViT models to diagnose COVID-19. A recent work by Shome *et al.* [238] attempts to mitigate this issue by combining three open-source COVID-19 datasets to create a large dataset comprising 30,000 images. Still, creating diverse and large COVID-19 datasets is challenging and requires significant effort from the medical community.

More attention must be given to design *interpretable* (to gain end-users trust) and *efficient* (for point-of-care testing) ViT models for COVID-19 diagnosis to make them a viable alternative of RT-PCR testing in the future.

We notice that most works have used the original ViT model [22] as a plug-and-play manner to boost the medical image classification performance. In this regard, we believe that integrating domain-specific context and accordingly designing architectural components and loss functions can enhance performance and provide

more insights in designing effective ViT-based classification models in the future.

Finally, let us highlight the exciting work of Matsoukas et al. [33] that, for the first time, demonstrates that ViTs pre-trained on ImageNet perform comparably to CNNs for the medical image classification task as shown in Table 5. This also raises an interesting question “Can ViT models pre-trained on medical imaging datasets perform better than ViT models pre-trained on ImageNet for medical image classification?”. A recent work by Xie et al. [277] attempts to answer this by pre-training the ViT on large-scale 2D and 3D medical images. On the medical image classification problem, their model obtains substantial performance gain over the ViT model pre-trained on ImageNet, indicating that this area is worth exploring further. A brief overview of ViT-based medical image classification approaches has been provided in Table 4.

5 MEDICAL IMAGE DETECTION

In medical image analysis, object detection refers to localization and identification of a region of interest (ROIs) such as lung nodules from X-ray images and is typically an essential aspect of diagnosis. However, it is one of the most time-consuming tasks for clinicians, thereby demanding the accurate computer-aided diagnosis (CAD) system to act as a second observer that may accelerate the process. Following the success of CNNs in medical image detection [278], [279], recently few attempts have been made to improve performance further using Transformer models. These approaches are mainly based on the detection transformer (DETR) framework [24].

Shen et al. [200] propose the first hybrid framework COTR, consisting of convolutional and transformer layers for end-to-end polyp detection. Specifically, the encoder of COTR contains six hybrid convolution-in-transformer layers to encode features. Whereas, the decoder consists of six transformer layers for object querying followed by a feed-forward network for object detection. COTR performs better than DETR on two different datasets ETIS-LARIB and CVC-ColonDB. The DETR model [24] is also adapted in other works [280], [281] for the end-to-end polyp detection [280], and detecting lymph nodes in T2 MRI scans for the assessment of lymphoproliferative diseases [281].

5.1 Discussion

Overall, the frequency of new Transformer-based approaches for the problem of medical image detection is lesser than those for the segmentation and classification. This is in contrast to the early years of CNN-based designs that were rapidly developed for the medical image detection, as indicated in Fig. 31. A recent work [282] shows that generic class-agnostic detection mechanism of multi-modal ViTs (like MDETR [283]) pre-trained on natural images-text pairs performs poorly on medical datasets. Therefore, investigating the performance of multi-modal ViTs by pre-training them on modality-specific medical imaging datasets is a promising future direction to explore. Furthermore, since the recent ViT-based methods yield competitive results on medical image detection problems, we expect to see more contributions in the near future.

6 MEDICAL IMAGE RECONSTRUCTION

The goal of medical image reconstruction is to obtain a clean image from a degraded input. For example, recovering a

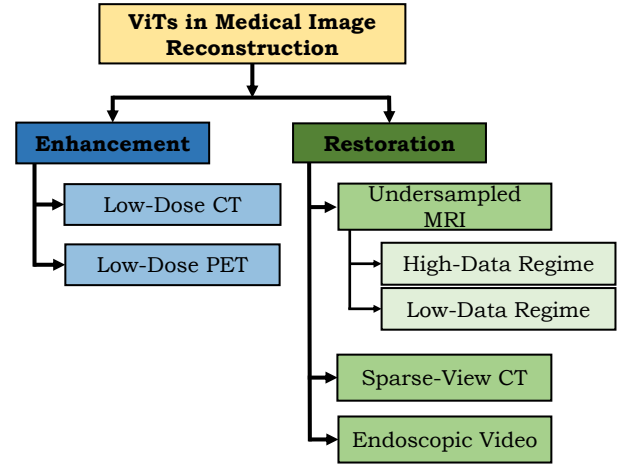


Figure 18: Taxonomy of ViT-based medical image reconstruction approaches.

high-resolution MRI image from its under-sampled version. It is a challenging task due to its ill-posed nature. Moreover, exact analytic inverse transforms in many practical medical imaging scenarios are unknown. Recently, ViTs have been shown to address these challenges effectively. We categorize the relevant works into *medical image enhancement* and *medical image restoration* areas, as depicted in Fig. 18.

6.1 Medical Image Enhancement

ViTs have achieved impressive success in the enhancement of medical images, mostly in the application of Low-Dose Computed Tomography (LDCT) [284], [285]. In LDCT, the X-ray dose is reduced to prevent patients from being exposed to high radiation. However, this reduction comes at the expense of CT image quality degradation and requires effective enhancement algorithms to improve the image quality and, subsequently, diagnostic accuracy.

6.1.1 LDCT Enhancement

Zhang et al. [37] propose a hybrid architecture TransCT that leverages the internal similarity of the LDCT images to effectively enhance them. TransCT first decomposes the LDCT image into high-frequency (HF) (containing noise) and low-frequency (LF) parts. Next it removes the noise from the HF part with the assistance of latent textures. To reconstruct the final high-quality LDCT images, TransCT further integrates features from the LF part to the output of the transformer decoder. Experiments on Mayo LDCT dataset [286] demonstrate the effectiveness of TransCT over CNN-based approaches. To perform LDCT image enhancement, Wang et al. [289] propose a convolution-free ViT-based encoder-decoder architecture TED-Net. It employs Token-to-token block [290] to enrich the image tokenization via a cascaded process. To refine contextual information, TED-Net introduces dilation and cyclic-shift blocks [125] in tokenization. TED-Net shows favorable performance on the Mayo Clinic LDCT dataset [286]. In another work, Luthra et al. [291] propose Eformer which is Transformer-based residual learning architecture for LDCT images denoising. To focus on edges, Eformer uses the power of Sobel-Feldman

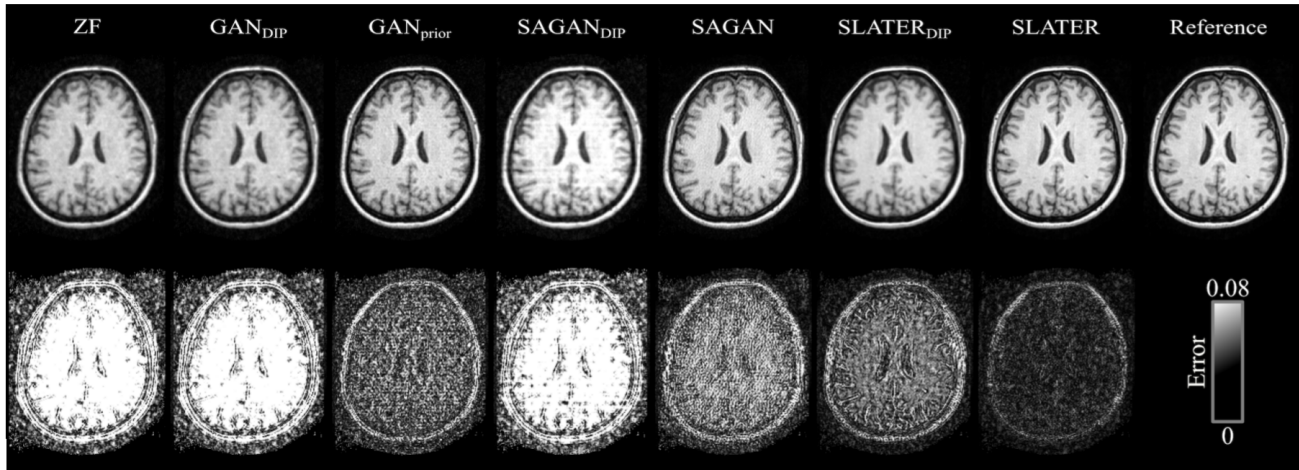


Figure 19: Unsupervised under-sampled MRI reconstruction results of CNN based and transformer based approaches. From left to right (top row): Fourier method (ZF) [287], GAN_{DIP} (unsupervised CNN), $\text{GAN}_{\text{prior}}$ [288] (unsupervised CNN with pre-training), Self-Attention GAN_{DIP} [46] (unsupervised CNN), Self-Attention GAN [288] (unsupervised CNN with pre-training), $\text{SLATER}_{\text{DIP}}$ [38] (unsupervised transformer), SLATER [38] (unsupervised transformer with pre-training), and reference image. Bottom row shows corresponding error maps. It can be seen that SLATER outperforms all other approaches in term of quality of reconstruction. Results taken from [38].

operator [292], [293] in the proposed edge enhancement block to boost denoising performance. Moreover, to handle the over-smoothness issue, the multi-scale perceptual loss [292] is used. Eformer achieves impressive image quality gains in terms PSNR, SSIM, and RMSE on the AAPM-Mayo Clinic dataset [286].

6.1.2 LDPET Enhancement

Like LDCT, Low-dose positron emission tomography (LDPET) images reduce the harmful radiation exposure of standard-dose PET (SDPET) at the expense of sacrificing diagnosis accuracy. To address this challenge, Luo *et al.* [294] propose an end-to-end generative adversarial network (GAN) based method integrated with Transformers, namely Transformer-GAN, to effectively reconstruct SDPET images from the corresponding LDPET images. Specifically, the generator of Transformer-GAN consists of a CNN-based encoder to learn compact feature representation, a transformer network to encode global context, and a CNN-based decoder to restore feature representation. They also introduce adversarial loss to obtain reliable and clinically acceptable images. Extensive experiments on their in-house collected clinical human brain PET dataset show the effectiveness of Transformer-GAN quantitatively and qualitatively.

6.2 Medical Image Restoration

Medical image restoration entails transforming signals collected by acquisition hardware (like MRI scanners) into interpretable images that can be used for diagnosis and treatment planning. Recently, ViT-based models have been proposed for multiple medical image restoration tasks, including undersampled MRI restoration, Sparse-View CT image reconstruction, and endoscopic video reconstruction. These models have pushed the boundaries of existing learning-based systems in terms of reconstruction accuracy. Next, we briefly highlight these approaches.

6.2.1 Undersampled MRI Reconstruction

Reducing the number of MRI measurements can result in faster scan times and a reduction in artifacts due to patients movement at the expense of aliasing artifacts in the image [295].

High-Data Regime Approaches. Approaches in this category assume the availability of large MRI training datasets to train the ViT model. Feng *et al.* [296] propose Transformer-based architecture, MTrans, for accelerated multi-modal MR imaging. The main component of MTrans is the cross attention module that extracts and fuses complementary features from the auxiliary modality to the target modality. Experiments on fastMRI and uiMRI datasets for reconstruction and super-resolution tasks show that MTrans achieve good performance gains over previous methods. However, MTrans requires separate training for MR reconstruction and super-resolution tasks. To jointly reconstruct and super-resolve MRI images, Feng *et al.* [297] propose Task-Transformer that leverages the power of multi-task learning to fuse complementary information between the reconstruction branch and the super-resolution branch. Experiments are performed on the public IXI and private MRI brain datasets. Similarly, Mahapatra *et al.* [298] propose a hybrid architecture to super-resolve MRI images by exploiting the complementary advantages of both CNNs and ViTs. They also propose novel loss functions [299] to preserve semantic and structural information in the super-resolved images.

Low-Data Regime Approaches. One drawback of the aforementioned approaches is the requirement of a massive paired dataset of undersampled and corresponding fully sampled MRI acquisitions to train ViT models. To alleviate the data requirement issue, Korkmaz *et al.* [38], [300] propose a zero-shot framework, SLATER, that leverages prior induced by randomly initialized neural networks [46], [47] for unsupervised MR image reconstruction. Specifically, during inference, SLATER inverts its transformer-based




Figure 20: **Top row:** Example reconstructions of models trained on 100 images (low-data regime) of fastMRI dataset. It can be seen that ViT model pre-trained on ImageNet produce sharp results as compared to recent CNN-based model [311] and randomly initialized ViT model. **Bottom row:** Example reconstructions of Brain images by models pre-trained on ImageNet and fine-tuned on Knee MRI dataset. Results show that pre-trained ViT models are more robust to anatomical shifts. Figures adapted from [301].

generative model via iterative optimization over-network weights to minimize the error between the network output and the under-sampled multi-coil MRI acquisitions while satisfying the MRI forward model constraints. SLATER yields quality improvements on single and multi-coil MRI brain datasets over other unsupervised learning-based approaches as shown in Fig. 19. Similarly, Lin *et al.* [301] show that a ViT model pre-trained on ImageNet, when fine-tuned on only 100 fastMRI images, not only yields sharp reconstructions but is also more robust towards anatomy shifts compared to CNNs as shown in Fig. 20. Furthermore, their experiments indicate that ViT benefits from higher throughput and less memory consumption than the U-Net baseline.

6.2.2 Sparse-View CT Reconstruction

Sparse-view CT [312] can effectively reduce the effective radiation dose by acquiring fewer projections. However, a decrease in the number of projections demands sophisticated image processing algorithms to achieve high-quality image reconstruction [313]. Wang *et al.* [305] present a hybrid CNN-Transformer, named Dual-Domain Transformer (DuDoTrans), by considering the global nature of sinogram’s sampling process to better restore high-quality images. In the first step, DuDoTrans reconstructs low-quality reconstructions of sinogram via filtered back projection step and learnable DuDo consistency layer. In the second step, a residual image reconstruction module performs enhancement to yield high-quality images. Experiments are performed on the NIH-AAPM dataset [286] to show generalizability, and robustness (against noise and artifacts) of DuDoTrans.

6.2.3 Endoscopic Video Reconstruction

Reconstructing surgical scenes from a stereoscopic video is challenging due to surgical tool occlusion and camera viewpoint changes. Long *et al.* [307] propose E-DSSR to reconstruct

surgical scenes from stereo endoscopic videos. Specifically, E-DSSR contains a lightweight stereo Transformer module to estimate depth images with high confidence and a segmentor network to accurately predict the surgical tool’s mask. Extensive experiments on Hamlyn Centre Endoscopic Video Dataset [308] and privately collected DaVinci robotic surgery dataset demonstrate the robustness of E-DSSR against abrupt camera movements and tissue deformations in real-time.

6.3 Discussion

In this section, we have reviewed about a dozen papers related to the applications of ViT models in medical image reconstruction, as shown in Table 6. Below, we highlight a few challenging problems and provide recent trends in the field.

Recently, an interesting work [301] has investigated the impact of pre-training ViT on the task of MRI image reconstruction. Their results indicate that pre-trained ViT yields sharp reconstructions and is robust towards anatomy shifts (see Fig. 20). The robustness of ViTs can be of particular relevance to the pathology image reconstruction as the range of pathology can vary significantly in the anatomy being imaged. Further, it raises an interesting question “Are ViTs pre-trained on medical image dataset able to provide any advantages in terms of reconstruction performance and robustness against anatomy shifts compared to their counterparts pre-trained on ImageNet”? Extensive and systematic experiments are required to answer this question. Another promising future direction is to investigate the impact on the performance of ViT pre-trained on one image modality (like CT) and fine-tuned on another modality (like MRI) for image reconstruction tasks.

We notice that most of the Transformer-based approach focus on MRI and CT image reconstruction tasks, and their applicability to other modalities are yet to be explored. In addition, proposed architectures are mostly generic and have not fully exploited the application-specific aspects. We believe that designing architectural components and formulating loss functions according to the task at hand can significantly boost performance.

We want to highlight one particular work that uses the Transformer-layer architecture to regularize the challenging problem of MRI image reconstruction from under-sampled measurements [38]. This work is inspired by the strong prior induced by the structure of untrained neural networks [46], [47]. These untrained network priors have recently garnered much attention from the medical image community as they do not need labeled training data. Considering advances in the untrained neural network area, we believe this direction requires further attention from medical imaging researchers in the context of Transformers.

We also observe that compared to the early years of CNNs (one paper from 2012 to 2015), Transformers have rapidly gained widespread attention in the medical image reconstruction community (more than a dozen papers in 2021), potentially due to the recent advancement in image-to-image translation frameworks.

7 MEDICAL IMAGE SYNTHESIS

In this section, we provide an overview of the applications of ViTs in medical image synthesis. Most of these approaches incorporate adversarial loss to synthesize realistic and high-quality medical images, albeit at the expense of training instability [314]. We have further classified these approaches into *intra-modality synthesis* and *inter-modality synthesis* due

Method	Highlights	Modality	Input Type	Datasets	Metric
TransCT [37]	Transformer for LDCT enhancement with high and low frequency decomposition.	CT	2D	NIH-AAPM [286]	RMSE, SSIM, VIF [302]
SLATER [38]	Transformer based approach for zero shot MRI image reconstruction.	MRI	3D	IXI [303] fastMRI [304]	PSNR, SSIM
TED-Net [289]	Pure transformer based encoder decoder dilation architecture for LDCT denoising.	CT	2D	NIH-AAPM [286]	RMSE, SSIM
Eformer [291]	Transformers based residual image denoising. Incorporate learnable Sobel filters for edge enhancement.	CT	2D	NIH-AAPM [286]	PSNR, SSIM, RMSE
Transformer-GAN [294]	End-to-end GAN-based method integrated with Transformers to enhance LDPET images.	PET	3D	Private	PSNR, SSIM, MSE
MTrans [296]	Leverage cross-attention module to fuse complementary features from the auxiliary modality to the target modality for fast multi-modal MRI image reconstruction.	MRI	2D	fastMRI [304] uiMRI (private)	PSNR, SSIM, NMSE
Task-Transformer [297]	Simultaneously, reconstruct and super-resolve MRI images via multi-task learning.	MRI	2D	IXI [303]	PSNR, SSIM, NMSE
Mahapatra <i>et al.</i> [298]	Hybrid architecture to super-resolve MRI images by exploiting the complementary advantages of CNNs and ViTs.	MRI	2D	fastMRI [304] IXI [303]	PSNR, SSIM, NMSE
Lin <i>et al.</i> [301]	ViT pretrained on ImageNet, when fine-tuned on only 100 fastMRI images, yields sharp reconstructions and is robust towards anatomy shifts.	MRI	2D	fastMRI [304]	SSIM
DuDoTrans [305]	A hybrid CNN-Transformer architecture that consider the global nature of sinogram's sampling process to restore high-quality CT images from sparse views.	CT	2D	NIH-AAPM [286]	PSNR, SSIM
MIST-Net [306]	Swin-transformer based projection and image domain two-stage framework to reconstruct high-quality CT images from sparse views.	CT	2D	NIH-AAPM [286]	PSNR, SSIM, RMSE
E-DSSR [307]	Leverage lightweight stereo Transformer module to estimate depth images with high confidence and a segmentor network to accurately predict the surgical tool's mask.	Endo.	2D	Hamlyn [308] DaVinci (private)	PSNR, SSIM
TransSMS [309]	ViT-based data consistency module to super-resolve magnetic particle imaging (MPI) system matrices for accelerated calibration.	MPI	2D, 3D	Open MPI [310] Private datasets	RMSE

Table 6: An overview of ViT-based approaches for medical image reconstruction.




Figure 21: Taxonomy of ViT-based medical image synthesis approaches.

to a different set of challenges in both categories, as shown in Fig. 21.

7.1 Intra-Modality Approaches

The goal of intra-modality synthesis is to generate higher-quality images from the relatively lower quality input images of the same modality. Next, we describe the details of ViT-based intra-modality medical image synthesis approaches.

7.1.1 Supervised Methods

Supervised image synthesis methods require paired source and target images to train ViT-based models. Paired data is difficult to obtain due to annotation cost and time constraints, thereby generally hindering the applicability of these models in medical imaging applications. Zhang *et al.* [315] focus on synthesizing infant brain structural MRIs (T1w and T2w scans) using both transformer and performer (simplified self-attention) layers [88]. Specifically, they design a novel multi-resolution pyramid-like U-Net framework, PTNet, utilizing performer encoder, performer decoder, and transformer bottleneck to synthesize high-quality infant MRI. They demonstrate the superiority of PTNet both qualitatively and quantitatively compared to pix2pix [316], and pix2pixHD [317] on large-scale infant MRI dataset [318]. Furthermore, in

addition to better synthesis quality, PTNet has a reasonable execution time of around 30 slices per second.

7.1.2 Semi-Supervised Methods

Semi-supervised approaches typically require small amounts of labeled data along with large unlabelled data to train models effectively. Kamran *et al.* [319] propose a multi-scale conditional generative adversarial network (GAN) [316] using ViT as a discriminator. They train their proposed model in a semi-supervised way to simultaneously synthesize Fluorescein Angiography (FA) images from fundus photographs and predict retinal degeneration. They use softmax activation after MLP head output and a categorical CE loss for classification. Besides adversarial loss, they also use MSE and perceptual losses to train their network. For ViT discriminator, they use embedding feature loss calculated using positional and patch features from the transformer encoder layers by successfully inserting the real and synthesized FA images. Their quantitative results in terms of Frechet inception Distance [320] and Kernel Inception Distance [321] demonstrate the superiority of their approach over baseline methods on diabetic retinopathy dataset provided by Hajeb *et al.* [322].

7.1.3 Unsupervised Methods

These approaches are particularly suitable for medical image synthesis tasks as they do not require paired training datasets. Recently, Ristea [323] proposed a cycle-consistent generative adversarial transformer (CyTran) to translate unpaired contrast CT scans to non-contrast CT scans and volumetric image registration of contrast CT scans to non-contrast CT scans. To handle high-resolution CT images, they propose hybrid convolution and multi-head attention-based architecture shown in Fig. 22. CyTran is unsupervised due to the integration of cyclic loss. Moreover, they introduce the

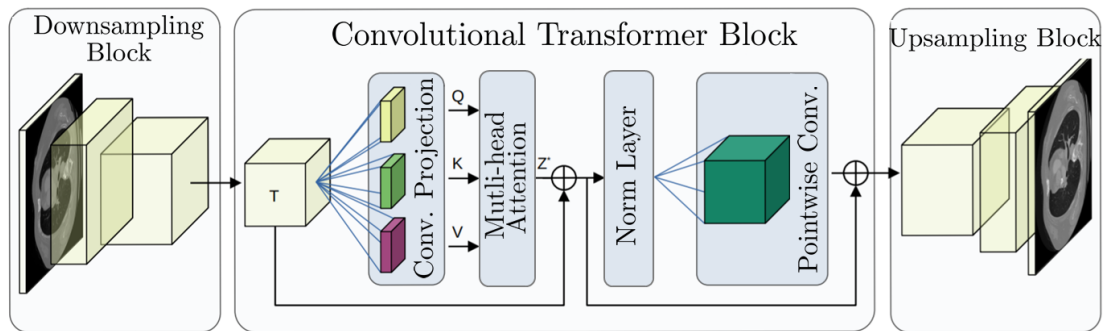


Figure 22: Hybrid convolutional-transformer network for CT image generation as proposed in [323]. It consists of downsampling convolutional layers to extract features from input images, a convolutional-transformer block comprising a multi-head self-attention mechanism, and an upsampling block to generate output images.




Figure 23: A pair of MRI (left) and CT (right) images of the same subject showing the significant appearance gap between the two modalities making medical image synthesis from MRI to CT a challenging task. Image is from [324].

Coltea-Lung-CT100W dataset formed of 100 3D anonymized triphasic lung CT scans of female patients.

7.2 Inter-Modality Approaches

The inter-modality approaches aim to synthesize targets to capture the useful structural information in the source images of different modalities. Examples include CT to MRI translation or vice-versa. Due to challenges associated with inter-modal translation, only supervised approaches have been explored.

Dalmaz *et al.* [325] introduce a novel synthesis approach, ResViT, for the multi-modal imaging based on a conditional deep adversarial network with ViT-based generator. Specifically, ResViT, employs convolutional and transformer branches within a residual bottleneck to preserve both local precision and contextual sensitivity along with the realism of adversarial learning. The bottleneck comprises novel aggregated residual transformer blocks to synergistically preserve local and global context, with a weight-sharing strategy to minimize model complexity. The effectiveness of ResViT model is demonstrated on two multi-contrast brain MRI datasets, BraTS [326]), and a multi-modal pelvic MRI-CT dataset [327].

7.3 Discussion

In this section, we have reviewed the applications of ViT models in medical image synthesis. Realistic synthesis of medical images

is particularly important as, in general, more than one imaging modality is involved in accurate clinical decision-making due to their complementary strengths. Therefore, in many practical applications, a certain modality is desired but infeasible to acquire due to cost and privacy issues. Recent transformer-based approaches can effectively circumvent these issues due to their ability to generate more realistic images than GAN-based methods.

Furthermore, most Transformer-based medical image synthesis approaches use the adversarial loss to generate realistic images. The adversarial loss can cause mode-collapse, and effective strategies must be employed to mitigate this issue [328].

Lastly, to the best of our knowledge, no work has been done using transformer-based models for inter-modality image synthesis approaches in an unsupervised setting. This can be due to the highly challenging nature of the problem (like CT and MRI images of the same subject have significantly different appearances, as shown in Fig. 23), thereby making it a promising direction to explore.

8 MEDICAL IMAGE REGISTRATION

Medical image registration aims to find dense per-voxel displacement and establish alignment between a pair of fixed and moving images. In medical imaging, registration may be necessary when analyzing a pair of images acquired at different times, from different viewpoints, or using different modalities (like MRI and CT) [75]. Accurate medical image registration is a challenging task due to difficulties in extracting discriminative features from multimodal medical images, complex motion, and lack of robust outlier rejection approaches [329]. In this section, we briefly highlight the applications of ViTs in medical image registration.

The first study to investigate the usage of transformers for self-supervised medical volumetric image registration has been proposed by Chen *et al.* [330]. Their model, ViT-V-Net, consists of a hybrid architecture composed of convolutional and transformer layers. Specifically, ViTs are applied to the high-level features of fixed and moving images extracted via a series of convolutional and max-pooling layers. The output from ViT is then reshaped and decoded using a V-Net style decoder [331]. To efficiently propagate the information, ViT-V-Net uses long skip connections between the encoder and decoder. The output of the ViT-V-Net decoder is a dense displacement field, which is fed to the

spatial transformer network for warping. Experiments on in-house MRI dataset show superiority of ViT-V-Net over other competing approaches in terms of Dice score. Chen *et al.* [332] further extends ViT-V-Net and propose TransMorph model for volumetric medical image registration. Particularly, TransMorph makes use of Swin Transformer in the encoder to capture the semantic correspondence between input fixed and moving images, followed by long skip connections-based convolutional decoder to predict dense displacement field. For uncertainty estimation, they also introduce Bayesian deep learning by applying variational inference on the parameters of the encoder in TransMorph. Extensive evaluation is performed to compare TransMorph with other approaches for the medical image registration task. Specifically, experiments on inter-patient brain MRI registration provided by John-Hopkin university and XCAT-to-CT registration demonstrate the superiority of TransMorph against twelve different hand-crafted, CNN-based, and transformer-based approaches. Similarly, Zhang *et al.* [333] present a novel dual transformer architecture (DTN) for volumetric diffeomorphic registration by effectively establishing correspondences between anatomical structures in an unsupervised manner. The DTN consists of two CNN-based 3D U-Net encoders to extract embeddings of separate and concatenated volumetric MRI images. To further refine and enhance the embeddings, they propose encoder-decoder-based dual transformers to encode the cross-volume dependencies. Given the enhanced embeddings, the CNN decoder infers the deformation fields. Qualitative and quantitative results in terms of Dice similarity coefficient and negative Jacobian determinant on OASIS dataset [334] of MRI scans demonstrate the effectiveness of their proposed architecture.

8.1 Discussion

Application of transformers to medical image registration problem is still at early stages, and it is difficult to draw any conclusion at this stage. However, seeing the rapid development of Transformer-based registration approaches in generic computer vision, we expect to see the same trend in this field in the near future.

9 CLINICAL REPORT GENERATION

Recently, immense progress has been made to automatically generate clinical reports from medical images using deep learning [335]–[338]. This automatic report generation process can help clinicians in accurate decision-making. However, generating reports (or captions) from the medical imaging data is challenging due to diversity in the reports of different radiologists, long sequence length (unlike natural image captions), and dataset bias (more normal data compared to abnormal). Moreover, an effective medical report generation model is expected to process two key attributes: (1) *language fluency* for human readability and (2) *clinical accuracy* to correctly identify the disease along with related symptoms. In this section, we briefly describe how transformer models help achieve these desired goals and effectively mitigate the aforementioned challenges associated with medical report generation. Specifically, these transformer-based approaches have achieved state-of-the-art performance both in terms of Natural Language Generation

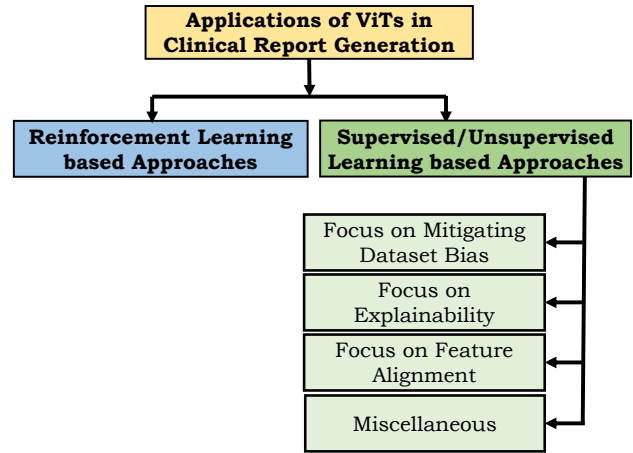


Figure 24: Taxonomy of applications of ViTs in clinical report generation.

(NLG) and Clinical Efficacy (CE) metrics. Also note that, unlike previous sections that mainly discuss ViTs, in this section, the focus is on the transformers as powerful language models to exploit the long-range dependencies for sentence generation. We have broadly categorized transformer-based clinical report generation approaches into reinforcement learning (RL) based and supervised/unsupervised learning methods, as shown in Fig. 24, due to differences in their underlying training mechanism.

9.1 Reinforcement Learning Based Approaches

RL-based medical report generation approaches can directly use the evaluation metrics of interest (like human evaluation, relevant medical terminologies, etc.) as rewards and update the model parameters via policy gradient. All approaches covered in this section use the self-critical RL [339] approach to train models, which is more suitable for the report generation task compared to the conventional RL.

One of the first attempts to integrate transformer in clinical report generation has been made by Xiong *et al.* [340]. They propose Reinforced-Transformer for Medical Image Captioning (RTMIC) that consists of a pre-trained DenseNet [236] to identify the region of interest from the input medical image, followed by a transformer-based encoder to extract visual features. These features are given as input to the captioning decoder to generate sentences. All these modules are updated via self-critical reinforcement learning method during training on IU Chest X-ray dataset [341]. Similarly, Miura *et al.* [342] show that the high accuracy of automatic radiology reports as measured by natural language generation metrics such as BLEU [343] and CIDEr [344] are often incomplete and inconsistent. To address these challenges, Miura *et al.* [342] propose a transformer-based model that directly optimizes the two newly proposed reward functions using self-critical RL. The first reward promotes the coverage of radiology domain entities with corresponding reference reports, and the second reward promotes the consistency of the generated reports with their descriptions in the reference reports. Further, they combine these reward functions with the semantic equivalence metric

of BERTScore [345] that results in generated reports with better performance in terms of clinical metrics.

Surgical Instructions Generation. Inspired by the success of transformers in medical report generation, Zhang *et al.* [346] propose a transformer model to generate instructions from the surgical scenes. Lack of a predefined template, as in the case of medical report generation, makes generation of surgical instructions a challenging task. To handle this challenge, Zhang *et al.* [346] have proposed an encoder-decoder based architecture back-boned by a transformer model. Specifically, their proposed architecture, optimized via self-critical reinforcement learning [339], effectively models the dependencies for visual features, textual features, and visual-textual relational features to accurately generate surgical reports on the DAISI dataset [347].

9.2 Supervised and Unsupervised Approaches

Supervised/unsupervised approaches use differentiable loss functions to train models for medical report generation and do not interact with the environment via an agent. We have categorized supervised/unsupervised approaches into methods that focus on dataset bias, explainability, feature alignment, and miscellaneous categories based on the challenges these approaches address.

9.2.1 Dataset Bias

Dataset bias is a common problem in medical report generation as there are far more sentences describing normalities than abnormalities. To mitigate this bias, Srinivasan [348] propose a hierarchical classification approach using a transformer as a decoder. Specifically, the transformer decoder leverage attention between and across features obtained from reports, images, and tags for effective report generation. The architecture consists of *Abnormality Detection Network* to classify normal and abnormal images, *Tag Classification Net* to generate tags against images, and *Report Generation Net* that takes image features and tags as inputs to generate final reports. Experiments on IU Chest X-ray dataset [341] demonstrate the effectiveness of the proposed architectural components. Similarly, Liu *et al.* [349] try to imitate the work of radiologists by distilling posterior and prior knowledge to generate accurate radiology reports. Their proposed architecture consists of three modules of Posterior Knowledge Explorer (PoKE), Prior Knowledge Explorer (PrKE), and Multidomain Knowledge Distiller (MKD). Specifically, PoKE identifies the abnormal area in the input images (mitigate image data bias), PrKE explores relevant prior information from the radiological reports and medical knowledge graph (mitigate textual data bias), and MKD (based on transformer decoder) distills the posterior and prior knowledge to generate radiology report. In another work, You *et al.* [350] propose AlignTransformer to generate a medical report from X-ray images. Specifically, AlignTransformer consists of two modules: align hierarchical attention and multi-grained transformer. Align hierarchical attention module helps to better locate the abnormal region in the input medical images. On the other hand, multi-grained transformer leverages multi-grained visual features using adaptive exploiting attention [351] to accurately generate long medical reports. AlignTransformer achieves favorable performance on IU-Xray [341] and MIMIC-CXR [352] datasets.




Figure 25: A chest X-ray image and its accompanying report, which includes findings and impressions, with aligned visual and textual components highlighted in different colors. Figure taken from [353].

Dataset	Number	Image Modality	View*	Length*	Report	
					Language	Cases
IU X-ray [341]	7,470	X-Ray	2	32.5	Eng.	2,955
MIMIC-CXR [352]	377,110	X-Ray	1	53.2	Eng.	276,778
PadChest [355]	160,868	X-Ray	2	-	Es	22,710
CX-CHR [356]	45,598	X-Ray	2	66.9	Zh	40,410
DIARETDB1 [357]	89	CFP	1	-	Eng	89
MESSIDOR [261]	1,200	CFP	2	-	Fr	587
FFA-IR [354]	1,048,584	FFA	87	91.2	Eng/Zh	10,790
COV-CTR [358]	728	CT-Scans	1	77.3	Eng/Zh	728
DEN [359]	15,709	CFP+FFA	1	7	Eng	-
STARE [360]	397	CFP+FFA	5	-	Eng	397

Table 7: Statistics of existing medical report generation datasets where * means average number and - for datasets that do not provide relevant information. Most of the transformers based models for medical report generation use Open-IU and MIMIC-CXR for evaluating results. Table adapted from [354].

9.2.2 Feature Alignment

Feature alignment based approaches mainly focus on the accurate alignment of encoded representation of the medical images and corresponding text, which is crucial for the interaction and generation across modalities (images and text here) and subsequently for accurate report generation, as indicated in Fig. 25. To align better, Chen *et al.* [353] propose a cross-modal memory network to augment the transformer-based encoder-decoder model for radiology report generation. They design a shared memory to facilitate the alignment between the features of medical images and texts. Experiments on IU-Xray [341] and MIMIC-CXR [352] datasets demonstrate that the proposed model can better align image and text features as compared to baseline methods. Similarly, building on the shared-memory work of Chen *et al.* [353], Yan *et al.* [361] introduce a weakly supervised contrastive objective to favor reports that are semantically close to the target, thereby producing more clinically accurate outputs. Similarly, Amjoud *et al.* [362] investigate the impact on the report generation performance by modifying different architectural components of the model proposed by Chen *et al.* [353] including replacing visual extractor and changing the number of layers in transformer-based decoder.

9.2.3 Explainable Models

Explainability in medical report generation is crucial to improve trustworthiness for deploying models in clinical settings and a mean for extracting bounding boxes for lesion localization. For model explainability, Hou *et al.* [363] employ attention to identify regions of interest in the input image and demonstrate where the model is focusing for the resulting

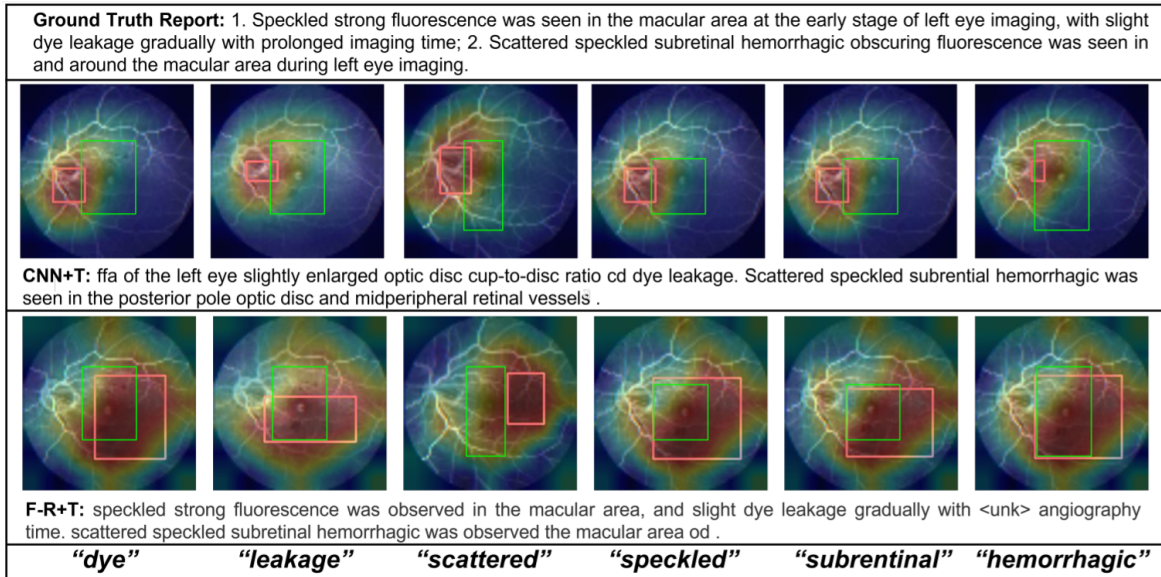


Figure 26: The depiction of lesion-image attention mapping areas and ground truth between CNN+Transformer and Faster-RCNN+Transformer samples, where green boxes represent the annotated region for each lesion word and red boxes represent the lesion-image attention mapping regions. Image taken from [354].

text. This attention mechanism increases the explainability of black-box models used in clinical settings and provides a method for extracting bounding boxes for disease localization. Specifically, they propose RATCHET transformer model to generate reports by using DenseNet-101 [236] as an image feature extractor. RATCHET consists of a transformer-based RNN-Decoder for generating chest radiograph reports. They assess the model’s natural language skills and the medical correctness of generated reports. Similarly, despite the immense interest of AI and clinical medicine researchers in the automatic report generation area, benchmark datasets are scarce, and the field lacks reliable evaluation metrics. To address these challenges, Li *et al.* [354] introduce a large-scale Fundus fluorescence in Angiography images and reports dataset containing 10,790 reports describing 1,048,584 images with explainable annotations as shown in Fig. 26. The dataset comes with annotated Chinese reports and corresponding translated English reports. Further, they introduce nine reliable metrics based on human evaluation criteria.

9.2.4 Miscellaneous

In this section, we highlight several approaches that try to improve different aspects of clinical report generation from medical images. Examples include a memory-driven transformer to capture similar patterns in reports, uncertainty quantification for reliable report generation, a curriculum learning-based method, and an unsupervised approach to avoid paired training datasets.

Chen *et al.* [364] propose a **memory-driven transformer to exploit similar patterns** in the radiology image reports. Specifically, they add a module to each layer of transformer-based decoder by optimizing the original layer normalization with a novel memory-driven conditional layer normalization. Extensive experiments on IU Chest X-ray [341] and MIMIC-CXR [352] datasets demonstrate the superiority of their approach both in terms of Natural Language Generation

(NLG) and Clinical Efficacy (CE) metrics. Similarly, Lovelace *et al.* [369] also leverage the transformer-based encoder and decoder for accurate medical report generation on MIMIC-CXR dataset [352]. To **emphasis on clinically relevant report generation**, they design a method to differentiate clinical information from generated reports, which they use to refine the model for clinical coherence. In another work, Alfarghaly *et al.* [372] present a pre-trained transformer-based model to generate a medical report from images. Specifically, the encoder consists of a pre-retrained CheXNet model that can generate semantic features from the input medical images. These semantic features are used to **condition GPT2 decoder** [373], [374] to generate accurate medical reports. Similarly, to judge the reliability of the automatic medical report generating model, uncertainty quantification is the key indicator. To incorporate this measure, Wang *et al.* [366] propose **transformer-based confidence guided framework** to quantify both visual and textual uncertainty. These uncertainties are subsequently used to construct an uncertainty-weighted loss to reduce misjudgment risk and improve the overall performance of the generated report. In other work, Nguyen *et al.* [367] propose differentiable end-to-end framework that consists of **transformer as generator for report generation**. Specifically, their proposed framework has three complementary modules: a *classifier* to learn the representation of disease features, a transformer-based *generator* model to generate the medical report, and *interpreter* to make the generated report consistent with the classifier output. They demonstrate the effectiveness of proposed components on IU-Xray [341] and MIMIC-CXR [352] datasets. Inspired by **curriculum learning** [375], Nooralahzadeh *et al.* [365] present a two-stage transformer architecture to progressively generate medical reports. Their progressive approach shows better performance over single-stage baselines in generating full-radiology reports. In another study, Pahwa *et al.* [376] investigate the **impact of visual feature extractor** model

DATASET	MODEL	YEAR	NLG METRICS						CE METRICS		
			BL-1	BL-2	BL-3	BL-4	MTR	RG-L	P	R	F1
IU X-RAY [341]	RTMIC [340]	2019	0.350	0.234	0.143	0.096	-	-	-	-	-
	HRG TRANSFORMER	2019	0.464	0.301	0.212	0.158	-	-	-	-	-
	KERP [356]	2019	0.482	0.325	0.226	0.162	-	0.339	-	-	-
	HIERARCHICAL TRANSFORMER [348]	2020	0.464	0.301	0.212	0.158	-	-	-	-	-
	MEMORY TRANSFORMER [364]	2020	0.470	0.304	0.219	0.165	0.187	0.371	-	-	-
	\mathcal{M}^2 TR. [365]	2021	0.475	0.301	0.228	0.180	0.169	0.373	-	-	-
	\mathcal{M}^2 TR. PROG. [365]	2021	0.486	0.317	0.232	0.173	0.192	0.390	-	-	-
	WANG ET AL. [366]	2021	0.481	0.309	0.223	0.169	0.193	0.365	-	-	-
	NGUYEN ET AL [367]	2021	0.515	0.378	0.293	0.235	0.219	0.436	-	-	-
	PPKED [349]	2021	0.483	0.315	0.224	0.168	0.190	0.376	-	-	-
	ALIGN TRANSFORMER [350]	2021	0.484	0.313	0.225	0.173	0.204	0.379	-	-	-
	KGAE UNSUPERVISED [368]	2021	0.417	0.263	0.181	0.126	0.149	0.318	-	-	-
	KGAE SEMI-SUPERVISED [368]	2021	0.497	0.320	0.232	0.171	0.189	0.379	-	-	-
	KGAE SUPERVISED [368]	2021	0.512	0.327	0.240	0.179	0.195	0.383	-	-	-
MIMIC-CXR [352]	TRANSFORMERS [14]	2017	0.409	0.268	0.191	0.144	0.157	0.318	-	-	-
	MEMORY TRANSFORMER [364]	2020	0.353	0.218	0.145	0.103	0.142	0.277	0.333	0.273	0.276
	CLINICAL TRANSFORMER [369]	2020	0.415	0.272	0.193	0.146	0.159	0.318	0.411	0.475	0.361
	\mathcal{M}^2 TR. [365]	2021	0.361	0.221	0.146	0.101	0.139	0.266	0.324	0.241	0.276
	\mathcal{M}^2 TR. PROG. [365]	2021	0.378	0.232	0.154	0.107	0.145	0.272	0.240	0.428	0.308
	PPKED [349]	2021	0.360	0.224	0.149	0.106	0.149	0.284	-	-	-
	ALIGN TRANSFORMER [350]	2021	0.378	0.235	0.156	0.112	0.158	0.283	-	-	-
	NGUYEN ET AL [367]	2021	0.495	0.360	0.278	0.224	0.222	0.390	-	-	-
	MDT+WCL [361]	2021	0.373	-	-	0.107	0.144	0.274	0.384	0.274	0.294
	\mathcal{M}^2 TRANS (CE) [342]	2021	-	-	-	0.111	-	-	0.463	0.732	0.567
	\mathcal{M}^2 TRANS (EN) [342]	2021	-	-	-	0.114	-	-	0.503	0.651	0.567
	KGAE UNSUPERVISED [368]	2021	0.221	0.144	0.096	0.062	0.097	0.208	0.214	0.158	0.156
	KGAE SEMI-SUPERVISED [368]	2021	0.352	0.219	0.149	0.108	0.147	0.290	0.360	0.302	0.307
	KGAE SUPERVISED [368]	2021	0.369	0.231	0.156	0.118	0.153	0.295	0.389	0.362	0.355

Table 8: Quantitative comparison of transformer models for the task of clinical report generation in terms of Natural Language Generation (NLG) and Clinical Efficacy (CE) on two benchmark datasets. The NLG metrics include BLEU (BL) [343], METEOR (MTR) [370] and ROUGE-L (RG-L) [371] and CE metrics include precision, recall and F1 score.

on the performance of medical report generation. Based on insights, they propose a modified HRNet [377], MedSkip, to extract visual features for the subsequent processing by the transformer-based decoder to generate an accurate medical report. Similarly, Park *et al.* [378] investigate the **expressiveness of features** to discriminate between normal and abnormal images. They demonstrate the superiority of transformer-based decoder without global average pooling over hierarchical LSTM baseline. Existing transformer-based report generation models are mostly supervised and use paired image-report data during training. The paired data is difficult to obtain due to privacy and cost in the medical domain. To mitigate this issue, Liu *et al.* [368] propose a knowledge graph auto-encoder that works in the shared latent domain of images and reports to extract useful information in an **unsupervised way**. Specifically, they use attention in the encoder to extract the knowledge representation from the knowledge graph and use a three-layer transformer in the decoder to generate reports. Their proposed framework can also be used in a semi-supervised or supervised manner in addition to the unsupervised mode. Quantitative and qualitative results, as well as evaluation by radiologists, corroborate the effectiveness of their approach.

9.3 Discussion

In this section, we have provided a comprehensive overview of the transformer’s applications for clinical report generation from X-ray images. In contrast to previous sections that discuss applications of ViTs, this section focuses on transformers as powerful language models. It is also pertinent to note that even though multiple surveys exist covering the applications of deep learning in clinical report generation [335]–[338], to the best of our knowledge, **none of these have covered the applications**

of transformer models in the area despite having transformers’ phenomenal impact since their inception back in 2017. In this regard, we hope this section will serve as a valuable resource to the research community. Below, we briefly highlight a few challenges associated with transformer-based models for report generation and outline promising future directions to explore.

As we have seen, transformer-based report generation models mostly rely on natural language generation (NLG) evaluation metrics such as CIDEr and BLEU to assess performance. These NLG metrics often fail to represent clinical efficacy. One recent work by Miura *et al.* [342] addresses this issue by proposing two new reward functions for the transformer model in reinforcement learning framework to better capture disease and anatomical knowledge in the generated reports. Another work by Li *et al.* [354] introduces nine reliable human evaluation criteria to validate the generated reports. Despite these works, we believe that more attention from the research community is required to design reliable clinical evaluation metrics to facilitate the adoption of transformer-based medical report generation models in clinical settings.

All transformer-based approaches covered in this section use the X-ray modality for automatic report generation. Generating reports from other modalities like MRI or PET have their own challenges associated with them due to their specific nature and distinct characteristics. Further, few medical datasets like ROCO [379], PEIR Gross [380], and ImageCLEF [381] are available that consist of multiple modalities, different body parts, and corresponding captions. These datasets have the potential to become worthy benchmarks to gauge the performance of future multimodal (or unimodal like MRI, PET) transformer-based models for medical report generation. We believe that transformer-based models tailored to specific modalities to generate reports must be explored in the future with a focus on creating diverse and challenging datasets of other modalities. Details of a few existing

medical report generation datasets are given in Table 7.

Further, we would like to point out the interested researchers toward the recently explored surgical instructions generation work using transformers [346] that could have a huge impact on surgical robotics, a market that is expected to reach USD 22.27 Billion by 2028. However, only one dataset, DAISI [347], is available to evaluate models in this emerging area, demanding attention from the medical community to create diverse and more challenging datasets.

Moreover, datasets for medical report generation like IU X-Ray [341] does not contain any standard train-test split and most of the transformer-based approaches evaluate the performance on different tests data. In this regard, the results in Table 8 are not directly comparable, but they can provide an overall indication about the performance of the models. We think what seems to be missing is a set of standardized procedures for creating challenging and diverse clinical report generation datasets.

10 OTHER APPLICATIONS

In this section, we briefly highlight applications of Transformers in other medical imaging areas, including survival outcome prediction, visual question answering, and medical point cloud analysis. **Survival outcome prediction** is a challenging regression task that seeks to predict the relative risk of cancer death. Recently, transformer models have shown impressive success in predicting survival rates. Chen *et al.* [382] propose a Multimodal Co-Attention Transformer (MCAT) for the survival outcome prediction from whole-slide imaging (WSI) in pathology. MCAT learns a co-attention mapping between genomics and WSIs features to discover how histology features attend to genes while predicting patient survival outcomes. Extensive experiments on five cancer datasets demonstrate the superiority of MCAT compared to state-of-the-art CNN-based approaches. Similarly, Kipkogei *et al.* [383] propose a Transformer-based architecture, Clinical Transformer, to model the relation between clinical and molecular features to predict survival outcomes from cancerous lung dataset [384]. In another work, Eslami [385] propose **PubMedCLIP, a fine-tuned version of Contrastive Language-Image Pre-Training (CLIP)** [386] for the medical domain by training it on the image-caption pairs from PubMed articles. Extensive experiments show that PubMedCLIP outperforms the previous state-of-the-art by nearly 3%. Recently, Liu *et al.* [10] propose **3D Medical Point Transformer (3DMPT)** to analyse 3D medical data. 3DMPT is tested on 3D medical classification and part segmentation based tasks. Similarly, Malkiel *et al.* [387] propose a **Transformer-based architecture to analyse fMRI data**. They pre-train the model on 4D fMRI data in a self-supervised manner and fine-tune it on various downstream tasks, including age and gender prediction, as well as diagnosing Schizophrenia.

11 OPEN CHALLENGES AND FUTURE DIRECTIONS

We have reviewed the exciting applications of vision transformers in medical image analysis. Despite their impressive performance, there remain several open research questions.

In this section, we outline some of their limitations and high-light promising future research directions. Specifically, we will discuss the challenges of pre-training on large datasets (Sec. 11.1), interpretability of ViT-based medical imaging approaches (Sec. 11.2), robustness against adversarial attacks (Sec. 11.3), designing efficient ViT architectures for real-time medical applications (Sec. 11.4), challenges in deploying ViT-based models in distributed settings (Sec. 11.5), and domain adaptation (Sec. 11.6). Further, wherever possible, we refer interested researchers to relevant CNNs-based medical imaging resources (recent studies, datasets, software libraries, etc.) to explore previously untapped applications by ViT-based models in medical imaging like adversarial robustness.

11.1 Pre-training

Due to a lack of intrinsic inductive biases in modeling local visual features, ViTs need to figure out the image-specific concepts on their own via pre-training from large-scale training datasets [22]. This may be a barrier to their widespread application in medical imaging, where typically datasets are orders of magnitude smaller compared to natural image datasets due to cost, privacy concerns, and the rarity of certain diseases, thereby making ViTs difficult to train efficiently in the medical domain. Existing learning-based medical imaging approaches commonly rely on transferring learning via ImageNet pretraining, which may be sub-optimal due to drastically different image characteristics between medical and natural images. Recently, Matsoukas *et al.* [33] has studied the impact of pre-training on ViTs performance for image classification and segmentation via a careful set of extensive experiments on several medical imaging datasets. Below, we briefly highlight major findings of their work.

- CNNs outperform ViTs for the medical image classification task when initialized with random weights.
- CNNs and ViTs benefit significantly from ImageNet initialization for medical image classification. ViTs appear to benefit more from transfer learning, as they make up for the gap observed using random initialization, performing on par with their CNN counterparts.
- CNNs and ViTs perform better with self-supervised pre-training approaches like DINO [271] and BYOL [388]. ViTs appear to outperform CNNs in this setting for medical image classification by a small margin.

In short, although recent ViT-based data-efficient approaches like DeiT [389], Token-to-Token [290], transformer in transformer [216], etc., report encouraging results in the generic vision applications, the task of learning these transformer models tailored to domain-specific medical imaging applications in a data-efficient manner is challenging. Recently, Tang *et al.* [163] has made an attempt to handle this issue by investigating the effectiveness of self-supervised learning as a pre-training strategy on domain-specific medical images. Specifically, they propose 3D transformer-based hierarchical encoder, Swin UNETR, and after pre-training on 5,050 CT images, demonstrates its effectiveness via fine-tuning on the downstream task of medical image segmentation. The pre-training on the medical imaging dataset also reduces




Figure 27: Impact of pre-training ViT on domain specific medical imaging dataset. *First column:* By only using 10% labelled data, Swin UNETR pre-trained on CT images (blue) is able to achieve 10% improvement in dice score over Swin UNETR trained from scratch (orange). *Middle column:* Dice scores of recently proposed transformer models on the spleen segmentation task of MSD dataset. Swin UNETR achieves state-of-the-art performance due to pre-training on the domain specific (CT) medical dataset. *Last column:* Qualitative visualizations of Swin UNETR pre-trained on CT images on BTCV multi-organ segmentation challenge. It can be seen that Swin UNETR predictions are closer to the groundtruth compared to baseline UNETR approach. First and last columns are adapted from [163].

the annotation effort compared to training Swin UNETR from scratch with random initialization. This is shown in Fig. 27, where it can be seen that pre-trained Swin UNETR can achieve the same performance by using only 60% of data as achieved by random initialized Swin UNETR using 100% of labeled data. This results in 40% reduction of manual annotation effort. Furthermore, as shown in Fig. 27, fine-tuning pre-trained Swin UNETR on the downstream medical image segmentation achieves better quantitative and qualitative results as compared to randomly initialized UNETR. Despite these efforts, there still remain several open challenges like Swin UNETR pre-trained on CT dataset gives unsatisfactory performance when applied directly to other medical imaging modalities like MRI due to large domain gap between CT and MRI images. Furthermore, the effectiveness of Swin UNETR on other downstream medical imaging tasks like classification and detection requires further investigation. Moreover, recent works for CNNs have shown that self-supervised pre-training on both ImageNet and medical datasets can improve the generalization performance (for classification) of the model on distribution shifted medical dataset [390] as compared to pre-training on ImageNet only. We believe such studies for ViT-based models, along with multi-instance contrastive learning to leverage patient meta data [391], will provide further insights to the community. Similarly, combining the self-supervised and semi-supervised pre-training in the context of ViTs for medical imaging applications is also an interesting avenue to explore [392].

11.2 Interpretability

Although the success of transformers has been empirically established in an impressive number of medical imaging applications, it has so far eluded a satisfactory interpretation. In most medical imaging applications, ViT models have been deployed as block-boxes, thereby failing to provide insights and explain their learning behavior for making predictions. This black-box nature of ViTs has hindered their deployment in clinical practice since, in areas such as medical applications, it is imperative to identify the limitations and potential failure cases of designed systems, where interpretability plays a

fundamental role [393]. Although several explainable AI-based medical imaging systems have been developed to gain deeper insights into the working of CNNs models for clinical applications [394]–[396], however, the work is still in its infancy for ViT-based medical imaging applications. It is despite the inherent suitability of the self-attention mechanism to interpretability due to its ability to explicitly model interactions between every region in the image as shown in Fig. 28 [397]. Recent efforts for interpretable ViT-based medical imaging models leverage saliency-based approaches [234] and Grad-CAM based visualizations [229]. Despite these efforts, the development of interpretable and explainable ViT-based approaches, specifically tailored for life-critical medical imaging applications, is a challenging and open research problem. Furthermore, formalisms, challenges, definitions, and evaluation protocols regarding interpretable ViTs based medical imaging systems must also be addressed. We believe that progress in this direction would not only help physicians to decide whether they should follow and trust automatic ViT-based model decisions but could also facilitate the deployment of such systems from a legal perspective.

11.3 Adversarial Robustness

Advances in adversarial attacks have revealed the vulnerability of existing learning-based medical imaging systems against imperceptible perturbation in the input images [404]–[406]. Considering the vast amount of money that underpins the medical imaging sector, this inevitably poses a risk whereby potential attackers may seek to profit from manipulation against the healthcare system, as shown in Fig. 29. For example, an attacker might try to manipulate the examination reports of patients for insurance fraud or a false medical reimbursement claim, thereby raising safety concerns. Therefore, ensuring the robustness of ViTs against adversarial attacks in life-critical medical imaging applications is of paramount importance. Although rich literature exists related to the robustness of CNNs in the medical imaging domain, to the best of our knowledge, no such study exists for ViTs, making it an exciting as well

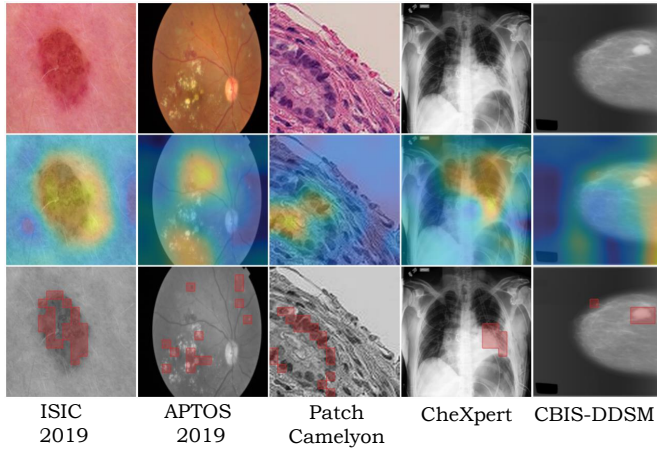


Figure 28: Saliency maps comparison for a CNN-based ResNet-50 [118] (second row) and ViT-based DEIT-S [389] (third row) on five medical imaging datasets for classification. Each column contains the original image (first row), a visualization of the ResNet-50 Grad-CAM saliency (second row), and a visualization of the DEIT-S’s attention map (third row), respectively. Note that the ViTs provide a clear, localized picture of attention compared to ResNet-50, thus giving insight into how the model makes decisions. Image taken from [33].

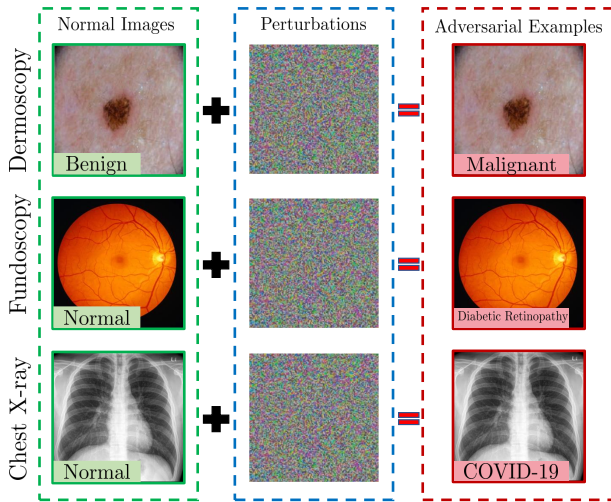


Figure 29: Examples of adversarial attacks to fool learning-based models trained on medical image datasets. Left: normal images, Middle: adversarial perturbations, Right: adversarial images. The bottom tag is the predicted class, and green/red indicates correct/wrong predictions.

challenging direction to explore. Recently, few attempts have been made to evaluate the robustness of ViTs to adversarial attacks for natural images [407]–[416]. The main conclusions of these attempts, ignoring their nuance difference, can be summarized as *ViTs are more robust to adversarial attacks than CNNs*. However, these robust ViT models cannot be directly deployed for medical imaging applications as the variety and type of patterns and textures in medical images differ significantly from the natural domain. Therefore, a principled approach to evaluate the robustness of ViTs against adversar-

Table 9: Description of datasets generally used in medical adversarial deep learning.

Dataset Name	Dataset Size	Modality
RSNA [264]	29,700	X-ray
JSRT [194]	247	X-ray
BraTS 2018 [398]	1689	MRI
BraTS 2019 [185]	1675	MRI
OASIS [399]	373-2168	MRI
HAM10000 [272]	10,000	Dermatoscopic
ISIC 18 [101]	3594	Dermatoscopic
LUNA 16 [400]	888	CT-Scans
NIH Chest X-ray [196]	112,000	X-ray
APTOS [255]	5590	Fundoscopy
Chest X-ray [401]	5856	X-ray
NSLT [402]	75,000	CT-Scans
Diabetic Retinopathy [403]	35,000	Fundoscopy

ial attacks in the medical imaging context, which builds the groundwork for resilience, could serve as a critical model to deploy these models in clinical settings. Furthermore, theoretical understanding to provide guarantees about the performance and robustness of ViTs, like CNNs [417], can be of significant interest to medical imaging researchers. In Table 9, we provide a description of datasets used in adversarial medical learning to evaluate the robustness of CNNs for interesting researchers to benchmark the robustness of ViT-based models.

11.4 ViTs for Medical Imaging on Edge Devices

Despite the tremendous success of ViTs in numerous medical imaging applications, the intensive requirements for memory and computation hamper their deployment on resource constraint edge devices [423], [424]. Due to recent advancements in edge computing, healthcare providers can process, store and analyze complex medical imaging data on-premises, speeding diagnosis, improving clinician workflows, enhancing patient privacy, and saving time—and potentially lives. These edge devices provide extremely fast and highly accurate processing of large amounts of medical imaging data, therefore demanding efficient hardware design to make ViT-based models suitable for edge computing-based medical imaging hardware. Recently few efforts have been made to compress transformer-based models by leveraging enhanced block-circulant matrix-based representation [425] and neural architecture search strategies [426]. Due to the exceptional performance of ViTs, we believe that there is a dire need for their domain-optimized architectural designs tailored for edge devices. It can have a tremendous impact on medical imaging-based health care applications where on-demand insights help teams make crucial and urgent decisions about patients.

11.5 Decentralized Medical Imaging Solutions using ViTs

Building robust deep learning-based medical imaging models highly depends on the amount and diversity of the training data. The training data required to train a reliable and robust model may not be available in a single institution due to strict privacy regulations, the low incidence rate of some pathologies, data-ownership concerns, and limited

Table 10: Description of some of the open-source federated learning and privacy-preserving frameworks, including some specifically developed for medical imaging applications.

Name	Framework	Description
TensorFlow Fed [418]	TensorFlow	Open-source framework for for machine learning and other computations on decentralized data.
CrypTen [419]	PyTorch	Framework to facilitate research in secure and privacy-preserving machine learning.
OpenMined [420]	TensorFlow, PyTorch, Keras	Open-source decentralised privacy preserving framework. Includes specialized tools like PySyft, PyGrid, and SyferText.
Opacus [421]	PyTorch	Enables training PyTorch models with differential privacy. Allows the client to online track the privacy budget.
Deepee [422]	PyTorch	Library for differentially private deep learning for medical imaging in PyTorch.
PriMIA [422]	PyTorch	Framework for end-to-end privacy-preserving decentralized deep learning for medical images.

numbers of patients. Federated Learning (FL) has been proposed to facilitate multi-hospital collaboration while obviating data transfer. Specifically, in FL, a shared model is built using distributed data from multiple devices where each device trains the model using its local data and then shares the model parameters with the central model without sharing its actual data. Although a plethora of approaches exists that address FL for CNNs based medical imaging applications, the work is still in its infancy for ViTs and requires further attention. Recently few research efforts have been made to exploit the inherent structure of ViT in distributed medical imaging applications. Park *et al.* [218] propose a Federated Split Task-Agnostic (FESTA) framework that integrates the power of Federated and Split Learning [219], [220] in utilizing ViT to simultaneously process multiple chest X-ray tasks, including diagnosing COVID-19 CXR images on a large corpus of decentralized data. Specifically, they split ViT into shared body and task-specific heads and demonstrate that ViT body with sufficient capacity can be shared across relevant tasks by leveraging the multitask-learning (MTL) [221] strategy. However, FESTA is just a proof-of-concept study, and its applicability in clinical trials requires further experimentation. Furthermore, challenges like privacy attacks and robustness against communication bottlenecks for ViT-based FL medical imaging systems require in-depth investigation. An interesting future direction is to explore recent privacy enhancement approaches like differential privacy [427] to prevent gradient inversion attacks [428] on FL-based medical imaging systems in the context of ViTs. In short, we believe that the successful implementation of distributed machine learning frameworks coupled with the strengths of ViTs could hold significant potential for enabling precision medicine at a large scale. This can lead to ViT models that yield unbiased decisions and are sensitive to rare diseases while respecting governance and privacy concerns. In Table 10, we highlight various tools and libraries that have been developed to implement distributed and secure deep learning. This can provide useful information for researchers who wish to rapidly prototype their ViT-based models for medical imaging in distributed settings.

11.6 Domain Adaptation and Out-of-Distribution Detection

Recent efforts for ViT-based medical imaging systems have primarily focused on improving the accuracy and gen-

erally lacking a principled mechanism to evaluate their generalization ability under different distribution/domain shifts. Recent studies have shown that test error generally increases in proportion to the distribution difference between training and test datasets, thereby making it a crucial issue to investigate in the context of ViTs. In medical imaging applications, these distribution shifts in data arise due to several factors that include: images acquired with a different device model at a different hospital, images of some unseen disease not in the training dataset, images that are incorrectly prepared, e.g., poor contrast, blurry images, etc. Extensive research exists on CNN-based out-of-distribution detection approaches in medical imaging [431]–[435]. Recently, few attempts have been made to show that large-scale pre-trained ViTs, due to their high-quality representations, can significantly improve the state-of-the-art on a range of out-of-distribution tasks across different data modalities [386], [430], [436]. However, investigation in these works has been mostly carried out on toy datasets such as CIFAR-10 and CIFAR-100, therefore not necessarily reflecting out-of-distribution detection performance on medical images with complex textures and patterns, high variance in feature scale (like in X-ray images), and local specific features. This demands further research to design ViT-based medical imaging systems that should be accurate for classes seen during training while providing calibrated estimates of uncertainty for abnormalities and unseen classes. We believe that research in this direction using techniques from transfer learning and domain adaptation will be of interest to the practitioners working in medical imaging based life-critical applications to envision potential practical deployment. In Fig. 30, we highlight the performance gain of ViTs as compared to CNNs for out of distribution detection to inspire medical imaging researchers who wish to explore this area. Another possible direction is to explore the recent advancements in continual learning [437] to effectively mitigate the issue of domains shift using ViTs. Few preliminary efforts have been made to explore this direction [438]; however, the work is still in its infancy and requires further attention from the community. Further, standardized and rigorous evaluation protocols also need to be established for domain adaptation in the medical imaging applications, similar to DOMAINBED [439] framework in the natural image domain. Such a framework will also help in advocating models reproducibility.

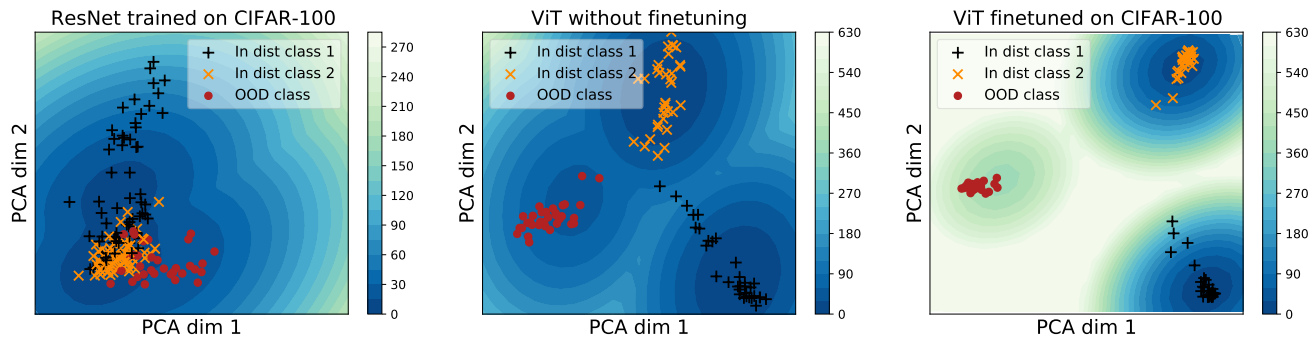


Figure 30: A 2D PCA projection of the space of embedding vectors for three models, having two in-distribution and one out-of-distribution class. The points are projections of embeddings of the categories of in-distribution classes (yellow and black) and out-of-distribution classes (red points). The color-coding shows the Mahalanobis outlier score [429]. ResNet-20 plot (left panel) leads to overlapping clusters indicating that classes are not well separated. ViT pre-trained on ImageNet-21k (middle panel) can distinguish classes from each other well but does not lead to well-separated outlier scores. ViT fine-tuned on the in-distribution dataset (right panel) is excellent at clustering embeddings based on classes and assigning high Mahalanobis distance to out-of-distribution inputs (red). Image courtesy [430].




Figure 31: ViT-based papers in medical imaging in 2021 (dark shade) vs pure CNN-based papers in medical imaging from 2012 to 2015 (light shade). It can be seen that ViTs have rapidly pervaded into almost all areas of medical imaging in a single year, with segmentation and classification being the most impacted areas. Statistics of CNNs-based papers are taken from [68].

12 DISCUSSION AND CONCLUSION

From the papers reviewed in this survey, it is evident that ViTs have pervaded every area of medical imaging (see Fig. 31). To keep pace with this rapid development, we recommend organizing the relevant workshops in top computer vision and medical imaging conferences and arranging special issues in prestigious journals to quickly disseminate the relevant research to the medical imaging

community.

In conclusion, we present the first comprehensive review of the applications of Transformers in medical imaging. We briefly cover the core concepts behind the success of Transformer models and then provide a comprehensive literature review of Transformers in a broad range of medical imaging tasks. Specifically, we survey the applications of Transformers in medical image segmentation, detection, classification, reconstruction, synthesis, registration, clinical report generation, and other tasks. In particular, for each of these applications, we develop taxonomy, identify application-specific challenges as well as give insights to solve them and specify recent trends. Despite their impressive performance, we anticipate there is still much exploration left to be done with Transformers in medical imaging, and we hope this survey provides a roadmap to researchers to progress this field further.

ACKNOWLEDGMENTS

The authors would like to thank Maryam Sultana (MBZ University of Artificial Intelligence) for her help with a few figures.

REFERENCES

- [1] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [2] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- [3] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [4] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *arXiv preprint arXiv:2201.03545*, 2022.
- [5] Paras Lakhani and Baskaran Sundaram. Deep learning at chest radiography: automated classification of pulmonary tuberculosis by using convolutional neural networks. *Radiology*, 284(2):574–582, 2017.
- [6] Jun Ki Min, Min Seob Kwak, and Jae Myung Cha. Overview of deep learning in gastrointestinal endoscopy. *Gut and liver*, 13(4):388, 2019.

- [7] Tobias Würfl, Florin C Ghesu, Vincent Christlein, and Andreas Maier. Deep learning computed tomography. In *International conference on medical image computing and computer-assisted intervention*, pages 432–440. Springer, 2016.
- [8] Michael M Lell and Marc Kachelrieß. Recent and upcoming technological developments in computed tomography: high speed, low dose, deep learning, multienergy. *Investigative radiology*, 55(1):8–19, 2020.
- [9] Azam Hamidinekoo, Erika Denton, Andrik Rampun, Kate Honnor, and Reyer Zwiggelaar. Deep learning in mammography and breast histology, an overview and future trends. *Medical image analysis*, 47:45–67, 2018.
- [10] Shengfeng Liu, Yi Wang, Xin Yang, Baiying Lei, Li Liu, Shawn Xi-ang Li, Dong Ni, and Tianfu Wang. Deep learning in medical ultrasound analysis: a review. *Engineering*, 5(2):261–275, 2019.
- [11] Alexander Selvikvåg Lundervold and Arvid Lundervold. An overview of deep learning in medical imaging focusing on mri. *Zeitschrift für Medizinische Physik*, 29(2):102–127, 2019.
- [12] Zeynettin Akkus, Alfiia Galimzianova, Assaf Hoogi, Daniel L Rubin, and Bradley J Erickson. Deep learning for brain mri segmentation: state of the art and future directions. *Journal of digital imaging*, 30(4):449–459, 2017.
- [13] Andrew J Reader, Guillaume Corda, Abolfazl Mehranian, Casper da Costa-Luis, Sam Ellis, and Julia A Schnabel. Deep learning for pet image reconstruction. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 5(1):1–25, 2020.
- [14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [15] 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.
- [16] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.
- [17] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [18] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. Disentangled non-local neural networks. In *European Conference on Computer Vision*, pages 191–207. Springer, 2020.
- [19] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jonathon Shlens. Stand-alone self-attention in vision models. *arXiv preprint arXiv:1906.05909*, 2019.
- [20] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3286–3295, 2019.
- [21] Ashish Vaswani, Prajit Ramachandran, Aravind Srinivas, Niki Parmar, Blake Hechtman, and Jonathon Shlens. Scaling local self-attention for parameter efficient visual backbones. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12894–12904, 2021.
- [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. *arXiv preprint arXiv:2010.11929*, 2020.
- [23] Sneha Chaudhari, Varun Mithal, Gungor Polatkan, and Rohan Ramanath. An attentive survey of attention models. *arXiv preprint arXiv:1904.02874*, 2019.
- [24] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.
- [25] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6881–6890, 2021.
- [26] Manoj Kumar, Dirk Weissenborn, and Nal Kalchbrenner. Colorization transformer. *arXiv preprint arXiv:2102.04432*, 2021.
- [27] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12299–12310, 2021.
- [28] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. *arXiv preprint arXiv:2103.15691*, 2021.
- [29] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *arXiv preprint arXiv:2105.10497*, 2021.
- [30] Eva Portelance, Michael C Frank, Dan Jurafsky, Alessandro Sordoni, and Romain Laroche. The emergence of the shape bias results from communicative efficiency. *arXiv preprint arXiv:2109.06232*, 2021.
- [31] Robert Geirhos, Kantharaju Narayanappa, Benjamin Mitzkus, Tizian Thieringer, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Partial success in closing the gap between human and machine vision. *arXiv preprint arXiv:2106.07411*, 2021.
- [32] Shikhar Tuli, Ishita Dasgupta, Erin Grant, and Thomas L Griffiths. Are convolutional neural networks or transformers more like human vision? *arXiv preprint arXiv:2105.07197*, 2021.
- [33] Christos Matsoukas, Johan Fredin Haslum, Magnus Söderberg, and Kevin Smith. Is it time to replace cnns with transformers for medical images? *arXiv preprint arXiv:2108.09038*, 2021.
- [34] Yunhe Gao, Mu Zhou, and Dimitris N Metaxas. Utnet: a hybrid transformer architecture for medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 61–71. Springer, 2021.
- [35] Ali Hatamizadeh, Dong Yang, Holger Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. *arXiv preprint arXiv:2103.10504*, 2021.
- [36] Mian Wu, Yinling Qian, Xiangyun Liao, Qiong Wang, and Pheng-Ann Heng. Hepatic vessel segmentation based on 3dswin-transformer with inductive biased multi-head self-attention. *arXiv preprint arXiv:2111.03368*, 2021.
- [37] Zhicheng Zhang, Lequan Yu, Xiaokun Liang, Wei Zhao, and Lei Xing. Transct: Dual-path transformer for low dose computed tomography. *arXiv preprint arXiv:2103.00634*, 2021.
- [38] Yilmaz Korkmaz, Salman UH Dar, Mahmut Yurt, Muzaffer Özbey, and Tolga Çukur. Unsupervised mri reconstruction via zero-shot learned adversarial transformers. *arXiv preprint arXiv:2105.08059*, 2021.
- [39] Yuanfeng Ji, Ruimao Zhang, Huijie Wang, Zhen Li, Lingyun Wu, Shaoting Zhang, and Ping Luo. Multi-compound transformer for accurate biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 326–336. Springer, 2021.
- [40] Paul Suetens. *Fundamentals of medical imaging*. Cambridge university press, 2017.
- [41] Xiaofeng Zhang, Nadine Smith, and Andrew Webb. Medical imaging. In *Biomedical Information Technology*, pages 3–27. Elsevier, 2008.
- [42] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.
- [43] Jun Xu, Lei Zhang, Wangmeng Zuo, David Zhang, and Xiangchu Feng. Patch group based nonlocal self-similarity prior learning for image denoising. In *Proceedings of the IEEE international conference on computer vision*, pages 244–252, 2015.
- [44] Yonina C Eldar and Gitta Kutyniok. *Compressed sensing: theory and applications*. Cambridge university press, 2012.
- [45] Hyeokho Choi, Justin Romberg, Richard Baraniuk, and Nick Kingsbury. Hidden markov tree modeling of complex wavelet transforms. In *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 00CH37100)*, volume 1, pages 133–136. IEEE, 2000.
- [46] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [47] Adnan Qayyum, Inaam Ilahi, Fahad Shamshad, Farid Boussaid, Mohammed Bennamoun, and Junaid Qadir. Untrained neural network priors for inverse imaging problems: A survey. 2021.
- [48] Adnan Qayyum, Waqas Sultani, Fahad Shamshad, Junaid Qadir, and Rashid Tufail. Single-shot retinal image enhancement using

- deep image priors. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 636–646. Springer, 2020.
- [49] KKD Ramesh, G Kiran Kumar, K Swapna, Debabrata Datta, and S Suman Rajest. A review of medical image segmentation algorithms. *EAI Endorsed Transactions on Pervasive Health and Technology*, 2021.
- [50] Saiprasad Ravishankar, Jong Chul Ye, and Jeffrey A Fessler. Image reconstruction: From sparsity to data-adaptive methods and machine learning. *Proceedings of the IEEE*, 108(1):86–109, 2019.
- [51] Eka Miranda, Mediana Aryuni, and E Irwansyah. A survey of medical image classification techniques. In *2016 International Conference on Information Management and Technology (ICIMTech)*, pages 56–61. IEEE, 2016.
- [52] Anand Jawdekar and Manish Dixit. A review of image enhancement techniques in medical imaging. *Machine Intelligence and Smart Systems*, pages 25–33, 2021.
- [53] Maximilian E Tschuchnig and Michael Gadermayr. Anomaly detection in medical imaging—a mini review. *arXiv preprint arXiv:2108.11986*, 2021.
- [54] Stefan Rüping et al. Learning interpretable models. 2006.
- [55] Muhammad Aurangzeb Ahmad, Carly Eckert, and Ankur Teredesai. Interpretable machine learning in healthcare. In *Proceedings of the 2018 ACM international conference on bioinformatics, computational biology, and health informatics*, pages 559–560, 2018.
- [56] S Kevin Zhou, Hayit Greenspan, Christos Davatzikos, James S Duncan, Bram Van Ginneken, Anant Madabhushi, Jerry L Prince, Daniel Rueckert, and Ronald M Summers. A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. *Proceedings of the IEEE*, 2021.
- [57] Chinmay Hegde. Algorithmic aspects of inverse problems using generative models. In *2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 166–172. IEEE, 2018.
- [58] Wenyi Lin, Kyle Hasenstab, Guilherme Moura Cunha, and Armin Schwartzman. Comparison of handcrafted features and convolutional neural networks for liver mr image adequacy assessment. *Scientific Reports*, 10(1):1–11, 2020.
- [59] Kjersti Engan, Sven Ole Aase, and J Hakon Husoy. Method of optimal directions for frame design. In *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No. 99CH36258)*, volume 5, pages 2443–2446. IEEE, 1999.
- [60] Michal Aharon, Michael Elad, and Alfred Bruckstein. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing*, 54(11):4311–4322, 2006.
- [61] Jian-Feng Cai, Hui Ji, Zuowei Shen, and Gui-Bo Ye. Data-driven tight frame construction and image denoising. *Applied and Computational Harmonic Analysis*, 37(1):89–105, 2014.
- [62] Xiaowei Zhou, Can Yang, Hongyu Zhao, and Weichuan Yu. Low-rank modeling and its applications in image analysis. *ACM Computing Surveys (CSUR)*, 47(2):1–33, 2014.
- [63] Jungha An, Mikael Rousson, and Chenyang Xu. γ -convergence approximation to piecewise smooth medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 495–502. Springer, 2007.
- [64] Eli Gibson, Wenqi Li, Carole Sudre, Lucas Fidon, Dzhoshkun I Shakir, Guotai Wang, Zach Eaton-Rosen, Robert Gray, Tom Doel, Yipeng Hu, et al. Niftynet: a deep-learning platform for medical imaging. *Computer methods and programs in biomedicine*, 158:113–122, 2018.
- [65] Fernando Pérez-García, Rachel Sparks, and Sebastien Ourselin. Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *Computer Methods and Programs in Biomedicine*, page 106236, 2021.
- [66] Andrew Beers, James Brown, Ken Chang, Katharina Hoebel, Jay Patel, K Ina Ly, Sara M Tolane, Priscilla Brastianos, Bruce Rosen, Elizabeth R Gerstner, et al. Deepneuro: an open-source deep learning toolbox for neuroimaging. *Neuroinformatics*, 19(1):127–140, 2021.
- [67] Xin Yi, Ekta Walia, and Paul Babyn. Generative adversarial network in medical imaging: A review. *Medical image analysis*, 58:101552, 2019.
- [68] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghahfouri, Jeroen AWM Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. A survey on deep learning in medical image analysis. *Medical image analysis*, 42:60–88, 2017.
- [69] Hayit Greenspan, Bram Van Ginneken, and Ronald M Summers. Guest editorial deep learning in medical imaging: Overview and future promise of an exciting new technique. *IEEE transactions on medical imaging*, 35(5):1153–1159, 2016.
- [70] S Kevin Zhou, Hayit Greenspan, and Dinggang Shen. *Deep learning for medical image analysis*. Academic Press, 2017.
- [71] Dinggang Shen, Guorong Wu, and Heung-Il Suk. Deep learning in medical image analysis. *Annual review of biomedical engineering*, 19:221–248, 2017.
- [72] Veronika Cheplygina, Marleen de Bruijne, and Josien PW Pluim. Not-so-supervised: a survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Medical image analysis*, 54:280–296, 2019.
- [73] Mohammad Hesam Hesamian, Wenjing Jia, Xiangjian He, and Paul Kennedy. Deep learning techniques for medical image segmentation: achievements and challenges. *Journal of digital imaging*, 32(4):582–596, 2019.
- [74] James S Duncan, Michael F Insana, and Nicholas Ayache. Biomedical imaging and analysis in the age of big data and deep learning [scanning the issue]. *Proceedings of the IEEE*, 108(1):3–10, 2019.
- [75] Grant Haskins, Uwe Kruger, and Pingkun Yan. Deep learning in medical image registration: a survey. *Machine Vision and Applications*, 31(1):1–18, 2020.
- [76] Nir Shlezinger, Jay Whang, Yonina C Eldar, and Alexandros G Dimakis. Model-based deep learning. *arXiv preprint arXiv:2012.08405*, 2020.
- [77] Vishal Monga, Yuelong Li, and Yonina C Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.
- [78] Gregory Ongie, Ajil Jalal, Christopher A Metzler, Richard G Baraniuk, Alexandros G Dimakis, and Rebecca Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020.
- [79] Rizwan Ahmad, Charles A Bouman, Gregory T Buzzard, Stanley Chan, Sizhuo Liu, Edward T Reehorst, and Philip Schniter. Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery. *IEEE signal processing magazine*, 37(1):105–116, 2020.
- [80] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [81] Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, et al. A survey on visual transformer. *arXiv preprint arXiv:2012.12556*, 2020.
- [82] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021.
- [83] Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*, 2020.
- [84] Tianyang Lin, Yuxin Wang, Xiangyang Liu, and Xipeng Qiu. A survey of transformers. *arXiv preprint arXiv:2106.04554*, 2021.
- [85] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *arXiv preprint arXiv:2106.02034*, 2021.
- [86] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pages 5156–5165. PMLR, 2020.
- [87] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nystromformer: A nystrom-based algorithm for approximating self-attention. *arXiv preprint arXiv:2102.03902*, 2021.
- [88] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020.

- [89] Noam Shazeer, Zhenzhong Lan, Youlong Cheng, Nan Ding, and Le Hou. Talking-heads attention. *arXiv preprint arXiv:2003.02436*, 2020.
- [90] Qiangguo Jin, Zhaopeng Meng, Changming Sun, Hui Cui, and Ran Su. Ra-unet: A hybrid deep attention-aware network to extract liver and tumor in ct scans. *Frontiers in Bioengineering and Biotechnology*, 8:1471, 2020.
- [91] Jo Schlemper, Ozan Oktay, Michiel Schaap, Mattias Heinrich, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention gated networks: Learning to leverage salient regions in medical images. *Medical image analysis*, 53:197–207, 2019.
- [92] Dhiraj Maji, Prarthana Sigedra, and Munendra Singh. Attention res-unet with guided decoder for semantic segmentation of brain tumors. *Biomedical Signal Processing and Control*, 71:103077, 2022.
- [93] Meng-Hao Guo, Tian-Xing Xu, Jiang-Jiang Liu, Zheng-Ning Liu, Peng-Tao Jiang, Tai-Jiang Mu, Song-Hai Zhang, Ralph R Martin, Ming-Ming Cheng, and Shi-Min Hu. Attention mechanisms in computer vision: A survey. *arXiv preprint arXiv:2111.07624*, 2021.
- [94] Danilo Avola, Luigi Cinque, Alessio Fagioli, Gianluca Foresti, and Alessio Mecca. Ultrasound medical imaging techniques: A survey. *ACM Computing Surveys (CSUR)*, 54(3):1–38, 2021.
- [95] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. *IEEE transactions on medical imaging*, 39(6):1856–1867, 2019.
- [96] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021.
- [97] Jiacheng Wang, Lan Wei, Liansheng Wang, Qichao Zhou, Lei Zhu, and Jing Qin. Boundary-aware transformers for skin lesion segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 206–216. Springer, 2021.
- [98] Yading Yuan. Automatic skin lesion segmentation with fully convolutional-deconvolutional networks. *arXiv preprint arXiv:1703.05165*, 2017.
- [99] David Gutman, Noel CF Codella, Emre Celebi, Brian Helba, Michael Marchetti, Nabin Mishra, and Allan Halpern. Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (isbi) 2016, hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1605.01397*, 2016.
- [100] Teresa Mendonça, Pedro M Ferreira, Jorge S Marques, André RS Marcal, and Jorge Rozeira. Ph 2-a dermoscopic image database for research and benchmarking. In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pages 5437–5440. IEEE, 2013.
- [101] Noel Codella, Veronica Rotemberg, Philipp Tschandl, M Emre Celebi, Stephen Dusza, David Gutman, Brian Helba, Aadi Kallou, Konstantinos Liopyris, Michael Marchetti, et al. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic). *arXiv preprint arXiv:1902.03368*, 2019.
- [102] Huisi Wu, Shihuai Chen, Guilian Chen, Wei Wang, Baiying Lei, and Zhenkun Wen. Fat-net: Feature adaptive transformers for automated skin lesion segmentation. *Medical Image Analysis*, page 102327, 2021.
- [103] Matt Berseth. Isic 2017-skin lesion analysis towards melanoma detection. *arXiv preprint arXiv:1703.00523*, 2017.
- [104] Hui Gao and Oksam Chae. Individual tooth segmentation from ct images using level set method with shape and intensity prior. *Pattern Recognition*, 43(7):2406–2417, 2010.
- [105] Yunxiang Li, Shuai Wang, Jun Wang, Guodong Zeng, Wenjun Liu, Qianni Zhang, Qun Jin, and Yaqi Wang. Gt u-net: A u-net like group transformer network for tooth root segmentation. In *International Workshop on Machine Learning in Medical Imaging*, pages 386–395. Springer, 2021.
- [106] Charles T Zahn and Ralph Z Roskies. Fourier descriptors for plane closed curves. *IEEE Transactions on computers*, 100(3):269–281, 1972.
- [107] Yunxiang Li, Guodong Zeng, Yifan Zhang, Jun Wang, Qun Jin, Lingling Sun, Qianni Zhang, Qisi Lian, Guiping Qian, Neng Xia, et al. Agmb-transformer: Anatomy-guided multi-branch transformer network for automated evaluation of root canal therapy. *IEEE Journal of Biomedical and Health Informatics*, 2021.
- [108] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.
- [109] Kaizhong Deng, Yanda Meng, Dongxu Gao, Joshua Bridge, Yaochun Shen, Gregory Lip, Yitian Zhao, and Yalin Zheng. Trans-bridge: A lightweight transformer for left ventricle segmentation in echocardiography. In *International Workshop on Advances in Simplifying Medical Ultrasound*, pages 63–72. Springer, 2021.
- [110] Qing-Long Zhang and Yu-Bin Yang. Sa-net: Shuffle attention for deep convolutional neural networks. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2235–2239. IEEE, 2021.
- [111] David Ouyang, Bryan He, Amirata Ghorbani, Neal Yuan, Joseph Ebinger, Curtis P Langlotz, Paul A Heidenreich, Robert A Harrington, David H Liang, Euan A Ashley, et al. Video-based ai for beat-to-beat assessment of cardiac function. *Nature*, 580(7802):252–256, 2020.
- [112] Yutong Xie, Jianpeng Zhang, Chunhua Shen, and Yong Xia. Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation. *arXiv preprint arXiv:2103.03024*, 2021.
- [113] Qian Wang, Li Sun, Yan Wang, Mei Zhou, Menghan Hu, Jiangang Chen, Ying Wen, and Qingli Li. Identification of melanoma from hyperspectral pathology image using 3d convolutional networks. *IEEE Transactions on Medical Imaging*, 40(1):218–227, 2020.
- [114] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [115] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.
- [116] Boxiang Yun, Yan Wang, Jieneng Chen, Huiyu Wang, Wei Shen, and Qingli Li. Spectr: Spectral transformer for hyperspectral pathology image segmentation. *arXiv preprint arXiv:2103.03604*, 2021.
- [117] Zhiqiang Shen, Hua Yang, Zhen Zhang, and Shaohua Zheng. Automated kidney tumor segmentation with convolution and transformer network. 2021.
- [118] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [119] The 2021 kidney and kidney tumor segmentation challenge. <https://kits21.kits-challenge.org/>. Accessed: 2022-01-20.
- [120] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017.
- [121] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [122] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [123] Hanchao Li, Pengfei Xiong, Jie An, and Lingxue Wang. Pyramid attention network for semantic segmentation. *arXiv preprint arXiv:1805.10180*, 2018.
- [124] Abhishek Chaurasia and Eugenio Culurciello. Linknet: Exploiting encoder representations for efficient semantic segmentation. In *2017 IEEE Visual Communications and Image Processing (VCIP)*, pages 1–4. IEEE, 2017.
- [125] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. *arXiv preprint arXiv:2105.05537*, 2021.
- [126] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.
- [127] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. *arXiv preprint arXiv:2105.05633*, 2021.

- [128] Jeya Maria Jose Valanarasu, Poojan Oza, Ilker Hacihaliloglu, and Vishal M Patel. Medical transformer: Gated axial-attention for medical image segmentation. *arXiv preprint arXiv:2102.10662*, 2021.
- [129] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [130] Yoo Jung Kim, Hyungjoon Jang, Kyoungbun Lee, Seongkeun Park, Sung-Gyu Min, Choyeon Hong, Jeong Hwan Park, Kanggeun Lee, Jisoo Kim, Wonjae Hong, et al. Paip 2019: Liver cancer segmentation challenge. *Medical Image Analysis*, 67:101854, 2021.
- [131] Cam Nguyen, Zuhayr Asad, and Yuankai Huo. Evaluating transformer based semantic segmentation networks for pathological image segmentation. *arXiv preprint arXiv:2108.11993*, 2021.
- [132] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020.
- [133] Tim Prangemeier, Christoph Reich, and Heinz Koepl. Attention-based transformers for instance segmentation of cells in microstructures. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 700–707. IEEE, 2020.
- [134] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [135] Bert Van den Bogerd, Nadia Zakaria, Bianca Adam, Steffi Matthyssen, Carina Koppen, and Sorchá Ní Dhubhghaill. Corneal endothelial cells over the past decade: are we missing the mark (er)? *Translational vision science & technology*, 8(6):13–13, 2019.
- [136] Yinglin Zhang, Risa Higashita, Huazhu Fu, Yanwu Xu, Yang Zhang, Haofeng Liu, Jian Zhang, and Jiang Liu. A multi-branch hybrid transformer network for corneal endothelial cell segmentation. *arXiv preprint arXiv:2106.07557*, 2021.
- [137] Alfredo Ruggeri, Fabio Scarpa, Massimo De Luca, Christian Meltendorf, and Jan Schroeter. A system for the automatic estimation of morphometric parameters of corneal endothelium in alizarine red-stained images. *British Journal of Ophthalmology*, 94(5):643–647, 2010.
- [138] Wenxuan Wang, Chen Chen, Meng Ding, Hong Yu, Sen Zha, and Jiangyun Li. Transbts: Multimodal brain tumor segmentation using transformer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 109–119. Springer, 2021.
- [139] Qiran Jia and Hai Shu. Bitr-unet: a cnn-transformer combined network for mri brain tumor segmentation. *arXiv preprint arXiv:2109.12271*, 2021.
- [140] Ujjwal Baid, Satyam Ghodasara, Suyash Mohan, Michel Bilello, Evan Calabrese, Errol Colak, Keyvan Farahani, Jayashree Kalpathy-Cramer, Felipe C Kitamura, Sarthak Pati, et al. The rsna-snr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314*, 2021.
- [141] Fabian Isensee, Jens Petersen, Andre Klein, David Zimmerer, Paul F Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Koehler, Tobias Norajitra, Sebastian Wirtk, et al. nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486*, 2018.
- [142] Louisa Lam and SY Suen. Application of majority voting to pattern recognition: an analysis of its behavior and performance. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 27(5):553–568, 1997.
- [143] Himashi Peiris, Munawar Hayat, Zhaolin Chen, Gary Egan, and Mehrtash Harandi. A volumetric transformer for accurate 3d tumor segmentation. *arXiv preprint arXiv:2111.13300*, 2021.
- [144] Hong-Yu Zhou, Jiansen Guo, Yinghao Zhang, Lequan Yu, Lian-sheng Wang, and Yizhou Yu. nnformer: Interleaved transformer for volumetric segmentation. *arXiv preprint arXiv:2109.03201*, 2021.
- [145] Ali Hatamizadeh, Vishwesh Nath, Yucheng Tang, Dong Yang, Holger Roth, and Daguang Xu. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. *arXiv preprint arXiv:2201.01266*, 2022.
- [146] Gonçalo M Correia, Vlad Niculae, and André FT Martins. Adaptively sparse transformers. *arXiv preprint arXiv:1909.00015*, 2019.
- [147] Qing Zhang, Qingli Li, Guanzhen Yu, Li Sun, Mei Zhou, and Junhao Chu. A multidimensional choledoch database and benchmarks for cholangiocarcinoma diagnosis. *IEEE access*, 7:149414–149421, 2019.
- [148] Qinghua Huang, Yaozhong Luo, and Qiangzhi Zhang. Breast ultrasound image segmentation: a survey. *International journal of computer assisted radiology and surgery*, 12(3):493–507, 2017.
- [149] Xiner Zhu, Haoji Hu, Hualiang Wang, Jincao Yao, Di Ou, Dong Xu, et al. Region aware transformer for automatic breast ultrasound tumor segmentation. 2021.
- [150] Yiyao Liu, Yi Yang, Wei Jiang, Tianfu Wang, and Baiying Lei. 3d deep attentive u-net with transformer for breast tumor segmentation from automated breast volume scanner. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 4011–4014. IEEE, 2021.
- [151] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- [152] Yang Lei, Yabo Fu, Tonghe Wang, Richard LJ Qiu, Walter J Curran, Tian Liu, and Xiaofeng Yang. Deep learning in multi-organ segmentation. *arXiv preprint arXiv:2001.10619*, 2020.
- [153] Davood Karimi, Serge Vasilychko, and Ali Gholipour. Convolution-free medical image segmentation using transformers. *arXiv preprint arXiv:2102.13645*, 2021.
- [154] Haoran Dou, Davood Karimi, Caitlin K Rollins, Cynthia M Ortinau, Lana Vasung, Clemente Velasco-Annis, Abdelhakim Ouallam, Xin Yang, Dong Ni, and Ali Gholipour. A deep attentive convolutional neural network for automatic cortical plate segmentation in fetal mri. *IEEE transactions on medical imaging*, 40(4):1123–1133, 2020.
- [155] Olivier Bernard, Alain Lalande, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE transactions on medical imaging*, 37(11):2514–2525, 2018.
- [156] Synapse multi-organ segmentation dataset. <https://www.synapse.org/#Synapse:syn3193805/wiki/217789>. Accessed: 2022-01-20.
- [157] Yundong Zhang, Huiye Liu, and Qiang Hu. Transfuse: Fusing transformers and cnns for medical image segmentation. *arXiv preprint arXiv:2102.08005*, 2021.
- [158] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [159] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *European Conference on Computer Vision*, pages 108–126. Springer, 2020.
- [160] Puyang Wang, Nick G Cuccolo, Rachana Tyagi, Ilker Hacihaliloglu, and Vishal M Patel. Automatic real-time cnn-based neonatal brain ventricles segmentation. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 716–719. IEEE, 2018.
- [161] 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. *Medical image analysis*, 35:489–502, 2017.
- [162] Neeraj Kumar, Ruchika Verma, Deepak Anand, Yanning Zhou, Omer Fahri Onder, Efstratios Tsougenis, Hao Chen, Pheng-Ann Heng, Jiahui Li, Zhiqiang Hu, et al. A multi-organ nucleus segmentation challenge. *IEEE transactions on medical imaging*, 39(5):1380–1391, 2019.
- [163] Yucheng Tang, Dong Yang, Wenqi Li, Holger Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis. *arXiv preprint arXiv:2111.14791*, 2021.
- [164] Ikboljon Sobirov, Otobek Nazarov, Hussain Alasmawi, and Mohammad Yaqub. Automatic segmentation of head and neck tumor: How powerful transformers are? *arXiv preprint arXiv:2201.06251*, 2022.
- [165] Yao Chang, Hu Menghan, Zhai Guangtao, and Zhang Xiao-Ping. Transclaw u-net: Claw u-net with transformers for medical image segmentation. *arXiv preprint arXiv:2107.05188*, 2021.
- [166] Chang Yao, Jingyu Tang, Menghan Hu, Yue Wu, Wenyi Guo, Qingli Li, and Xiao-Ping Zhang. Claw u-net: A unet-based

- network with deep feature concatenation for scleral blood vessel segmentation. *arXiv preprint arXiv:2010.10163*, 2020.
- [167] Ben Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. Levit: a vision transformer in convnet's clothing for faster inference. *arXiv preprint arXiv:2104.01136*, 2021.
- [168] Guoping Xu, Xingrong Wu, Xuan Zhang, and Xinwei He. LeViT-UNet: make faster encoders with transformer for medical image segmentation. *arXiv preprint arXiv:2107.08623*, 2021.
- [169] Bingzhi Chen, Yishu Liu, Zheng Zhang, Guangming Lu, and David Zhang. Transattunet: Multi-level attention-guided u-net with transformer for medical image segmentation. *arXiv preprint arXiv:2107.05274*, 2021.
- [170] Xiangyi Yan, Hao Tang, Shanlin Sun, Haoyu Ma, Deying Kong, and Xiaohui Xie. After-unet: Axial fusion transformer unet for medical image segmentation. *arXiv preprint arXiv:2110.10403*, 2021.
- [171] Amber L Simpson, Michela Antonelli, Spyridon Bakas, Michel Bilello, Keyvan Farahani, Bram Van Ginneken, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, et al. A large annotated medical image dataset for the development and evaluation of segmentation algorithms. *arXiv preprint arXiv:1902.09063*, 2019.
- [172] Xuming Chen, Shanlin Sun, Narisu Bai, Kun Han, Qianqian Liu, Shengyu Yao, Hao Tang, Chupeng Zhang, Zhipeng Lu, Qian Huang, et al. A deep learning-based auto-segmentation system for organs-at-risk on whole-body computed tomography images for radiation therapy. *Radiotherapy and Oncology*, 160:175–184, 2021.
- [173] Zoé Lambert, Caroline Petitjean, Bernard Dubray, and Su Kuan. Segthor: Segmentation of thoracic organs at risk in ct images. In 2020 Tenth International Conference on Image Processing Theory, Tools and Applications (IPTA), pages 1–6. IEEE, 2020.
- [174] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [175] Víctor M Campello, Polyxeni Gkontra, Cristian Izquierdo, Carlos Martín-Isla, Alireza Sojoudi, Peter M Full, Klaus Maier-Hein, Yao Zhang, Zhiqiang He, Jun Ma, et al. Multi-centre, multi-vendor and multi-disease cardiac segmentation: the m&ms challenge. *IEEE Transactions on Medical Imaging*, 40(12):3543–3554, 2021.
- [176] Ailiang Lin, Bingzhi Chen, Jiayu Xu, Zheng Zhang, and Guangming Lu. Ds-transunet: Dual swin transformer u-net for medical image segmentation. *arXiv preprint arXiv:2106.06716*, 2021.
- [177] Yijiang Li, Wentian Cai, Ying Gao, and Xiping Hu. More than encoder: Introducing transformer decoder to upsample. *arXiv preprint arXiv:2106.10637*, 2021.
- [178] Debesh Jha, Pia H Smedsrud, Michael A Riegler, Pål Halvorsen, Thomas de Lange, Dag Johansen, and Håvard D Johansen. Kvasir-seg: A segmented polyp dataset. In *International Conference on Multimedia Modeling*, pages 451–462. Springer, 2020.
- [179] Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Debora Gil, Cristina Rodríguez, and Fernando Vilarinho. Wm-dova maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians. *Computerized Medical Imaging and Graphics*, 43:99–111, 2015.
- [180] Nima Tajbakhsh, Suryakanth R Gurudu, and Jianming Liang. Automated polyp detection in colonoscopy videos using shape and context information. *IEEE transactions on medical imaging*, 35(2):630–644, 2015.
- [181] David Vázquez, Jorge Bernal, F Javier Sánchez, Gloria Fernández-Esparrach, Antonio M López, Adriana Romero, Michal Drozdal, and Aaron Courville. A benchmark for endoluminal scene segmentation of colonoscopy images. *Journal of healthcare engineering*, 2017, 2017.
- [182] Juan Silva, Aymeric Histace, Olivier Romain, Xavier Dray, and Bertrand Granado. Toward embedded detection of polyps in wce images for early diagnosis of colorectal cancer. *International journal of computer assisted radiology and surgery*, 9(2):283–293, 2014.
- [183] Zhuangzhuang Zhang, Baozhou Sun, and Weixiong Zhang. Pyramid medical transformer for medical image segmentation. *arXiv preprint arXiv:2104.14702*, 2021.
- [184] Neeraj Kumar, Ruchika Verma, Sanuj Sharma, Surabhi Bhargava, Abhishek Vahadane, and Amit Sethi. A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE transactions on medical imaging*, 36(7):1550–1560, 2017.
- [185] Multimodal brain tumor segmentation challenge 2019. <https://www.med.upenn.edu/cbica/brats-2019/>. Accessed: 2022-01-20.
- [186] Multimodal brain tumor segmentation challenge 2020. <https://www.med.upenn.edu/cbica/brats-2020/>. Accessed: 2022-01-20.
- [187] Olivier Petit, Nicolas Thome, Clement Rambour, Loic Themyr, Toby Collins, and Luc Soler. U-net transformer: self and cross attention for medical image segmentation. In *International Workshop on Machine Learning in Medical Imaging*, pages 267–276. Springer, 2021.
- [188] Kenneth Clark, Bruce Vendt, Kirk Smith, John Freymann, Justin Kirby, Paul Koppel, Stephen Moore, Stanley Phillips, David Maffitt, Michael Pringle, et al. The cancer imaging archive (cica): maintaining and operating a public information repository. *Journal of digital imaging*, 26(6):1045–1057, 2013.
- [189] Vincent Andrearczyk, Valentin Oreiller, Mario Jreige, Martin Vallières, Joel Castelli, Hesham Elhalawani, Sarah Boughdad, John O Prior, and Adrien Depeursinge. Overview of the hecktor challenge at miccai 2020: automatic head and neck tumor segmentation in pet/ct. In *3D Head and Neck Tumor Segmentation in PET/CT Challenge*, pages 1–21. Springer, 2020.
- [190] Shaohua Li, Xiuchao Sui, Xiangde Luo, Xinxing Xu, Yong Liu, and Rick Siow Mong Goh. Medical image segmentation using squeeze-and-expansion transformers. *arXiv preprint arXiv:2105.09511*, 2021.
- [191] José Ignacio Orlando, Huazhu Fu, João Barbosa Breda, Karel van Keer, Deepti R Bathula, Andrés Díaz-Pinto, Ruogu Fang, Pheng-Ann Heng, Jeyoung Kim, JoonHo Lee, et al. Refuge challenge: A unified framework for evaluating automated methods for glaucoma assessment from fundus photographs. *Medical image analysis*, 59:101570, 2020.
- [192] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Prnet: Parallel reverse attention network for polyp segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 263–273. Springer, 2020.
- [193] Juan C Caicedo, Allen Goodman, Kyle W Karhohs, Beth A Cimini, Jeanelle Ackerman, Marzieh Haghighi, CherKeng Heng, Tim Becker, Minh Doan, Claire McQuin, et al. Nucleus segmentation across imaging experiments: the 2018 data science bowl. *Nature methods*, 16(12):1247–1253, 2019.
- [194] Junji Shiraishi, Shigehiko Katsuragawa, Junpei Ikezoe, Tsuneo Matsumoto, Takeshi Kobayashi, Ken-ichi Komatsu, Mitate Matsui, Hiroshi Fujita, Yoshie Kodera, and Kunio Doi. Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules. *American Journal of Roentgenology*, 174(1):71–74, 2000.
- [195] Stefan Jaeger, Sema Candemir, Sameer Antani, Yi-Xiang J Wang, Pu-Xuan Lu, and George Thoma. Two public chest x-ray datasets for computer-aided screening of pulmonary diseases. *Quantitative imaging in medicine and surgery*, 4(6):475, 2014.
- [196] You-Bao Tang, Yu-Xing Tang, Jing Xiao, and Ronald M Summers. Xlsor: A robust and accurate lung segmentor on chest x-rays using criss-cross attention and customized radiorealistc abnormalities generation. In *International Conference on Medical Imaging with Deep Learning*, pages 457–467. PMLR, 2019.
- [197] Xin He, Shihao Wang, Shaohuai Shi, Xiaowen Chu, Jiangping Tang, Xin Liu, Chenggang Yan, Jiyong Zhang, and Guiguang Ding. Benchmarking deep learning models and automated model design for covid-19 detection with chest ct scans. *medRxiv*, 2020.
- [198] Bo Dong, Wenhui Wang, Deng-Ping Fan, Jinpeng Li, Huazhu Fu, and Ling Shao. Polyp-pvt: Polyp segmentation with pyramid vision transformers. *arXiv preprint arXiv:2108.06932*, 2021.
- [199] Xiaohong Huang, Zhifang Deng, Dandan Li, and Xueguang Yuan. Missformer: An effective medical image segmentation transformer. *arXiv preprint arXiv:2109.07162*, 2021.
- [200] Zhiqiang Shen, Chaonan Lin, and Shaohua Zheng. Cotr: Convolution in transformer network for end to end polyp detection. *arXiv preprint arXiv:2105.10925*, 2021.
- [201] Golnaz Ghiasi and Charles C Fowlkes. Laplacian pyramid reconstruction and refinement for semantic segmentation. In *European conference on computer vision*, pages 519–534. Springer, 2016.
- [202] Changqian Yu, Jingbo Wang, Changxin Gao, Gang Yu, Chunhua Shen, and Nong Sang. Context prior for scene segmentation. In

- Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12416–12425, 2020.
- [203] B Landman, Z Xu, J Eugenio Igelsias, M Styner, T Langerak, and A Klein. Miccai multi-atlas labeling beyond the cranial vault-workshop and challenge. In *Proc. MICCAI: Multi-Atlas Labeling Beyond Cranial Vault-Workshop Challenge*, 2015.
- [204] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017.
- [205] Tao Ai, Zhenlu Yang, Hongyan Hou, Chenao Zhan, Chong Chen, Wenzhi Lv, Qian Tao, Ziyong Sun, and Liming Xia. Correlation of chest ct and rt-pcr testing for coronavirus disease 2019 (covid-19) in china: a report of 1014 cases. *Radiology*, 296(2):E32–E40, 2020.
- [206] Yicheng Fang, Huangqi Zhang, Jicheng Xie, Minjie Lin, Lingjun Ying, Peipei Pang, and Wenbin Ji. Sensitivity of chest ct for covid-19: comparison to rt-pcr. *Radiology*, 296(2):E115–E117, 2020.
- [207] Dandan Chen, Xinqing Jiang, Yong Hong, Zhihui Wen, Shuquan Wei, Guangming Peng, and Xinhua Wei. Can chest ct features distinguish patients with negative from those with positive initial rt-pcr results for coronavirus disease (covid-19)? *American Journal of Roentgenology*, 216(1):66–70, 2021.
- [208] Shehan Perera, Srikar Adhikari, and Alper Yilmaz. Pocformer: A lightweight transformer architecture for detection of covid-19 using point of care ultrasound. *arXiv preprint arXiv:2105.09913*, 2021.
- [209] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [210] Jannis Born, Gabriel Brändle, Manuel Cossio, Marion Disdier, Julie Goulet, Jérémie Roulin, and Nina Wiedemann. Pocovid-net: automatic detection of covid-19 from a new lung ultrasound imaging dataset (pocus). *arXiv preprint arXiv:2004.12084*, 2020.
- [211] Joseph Paul Cohen, Paul Morrison, Lan Dao, Karsten Roth, Tim Q Duong, and Marzyeh Ghassemi. Covid-19 image data collection: Prospective predictions are the future. *arXiv preprint arXiv:2006.11988*, 2020.
- [212] Chengeng Liu and Qingshan Yin. Automatic diagnosis of covid-19 using a tailored transformer-like network. In *Journal of Physics: Conference Series*, volume 2010, page 012175. IOP Publishing, 2021.
- [213] Li Yuan, Qibin Hou, Zihang Jiang, Jiashi Feng, and Shuicheng Yan. Volc: Vision outlooker for visual recognition. *arXiv preprint arXiv:2106.13112*, 2021.
- [214] Muhammad EH Chowdhury, Tawsifur Rahman, Amith Khandakar, Rashid Mazhar, Muhammad Abdul Kadir, Zaid Bin Mahbub, Khandakar Reajul Islam, Muhammad Salman Khan, Atif Iqbal, Nasser Al Emadi, et al. Can ai help in screening viral and covid-19 pneumonia? *IEEE Access*, 8:132665–132676, 2020.
- [215] Juntao Jiang and Shuyi Lin. Covid-19 detection in chest x-ray images using swin-transformer and transformer in transformer. *arXiv preprint arXiv:2110.08427*, 2021.
- [216] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *arXiv preprint arXiv:2103.00112*, 2021.
- [217] Qi Dou, Tiffany Y So, Meirui Jiang, Quande Liu, Varut Vardhanabhuti, Georgios Kaissis, Zeju Li, Weixin Si, Heather HC Lee, Kevin Yu, et al. Federated deep learning for detecting covid-19 lung abnormalities in ct: a privacy-preserving multinational validation study. *NPJ digital medicine*, 4(1):1–11, 2021.
- [218] Sangjoon Park, Gwanghyun Kim, Jeongsol Kim, Boah Kim, and Jong Chul Ye. Federated split vision transformer for covid-19cxr diagnosis using task-agnostic training. *arXiv preprint arXiv:2111.01338*, 2021.
- [219] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- [220] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split learning for health: Distributed deep learning without sharing raw patient data. *arXiv preprint arXiv:1812.00564*, 2018.
- [221] Rich Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.
- [222] Thomas C Kwee and Robert M Kwee. Chest ct in covid-19: what the radiologist needs to know. *RadioGraphics*, 40(7):1848–1865, 2020.
- [223] Chih-Chung Hsu, Guan-Lin Chen, and Mei-Hsuan Wu. Visual transformer with statistical test for covid-19 classification. *arXiv preprint arXiv:2107.05334*, 2021.
- [224] Robert F Woolson. Wilcoxon signed-rank test. *Wiley encyclopedia of clinical trials*, pages 1–3, 2007.
- [225] Lei Zhang and Yan Wen. A transformer-based framework for automatic covid19 diagnosis in chest cts. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 513–518, 2021.
- [226] Lei Zhang and Yan Wen. Mia-cov19d: A transformer-based framework for covid19 classification in chest cts.
- [227] Dimitrios Kollias, Anastasios Arsenos, Levon Soukissian, and Stefanos Kollias. Mia-cov19d: Covid-19 detection through 3-d chest ct image analysis. *arXiv preprint arXiv:2106.07524*, 2021.
- [228] Runmin Cong, Jianjun Lei, Huazhu Fu, Ming-Ming Cheng, Weisi Lin, and Qingming Huang. Review of visual saliency detection with comprehensive information. *IEEE Transactions on circuits and Systems for Video Technology*, 29(10):2941–2959, 2018.
- [229] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [230] Arnab Kumar Mondal, Arnab Bhattacharjee, Parag Singla, and Prathosh AP. xvitos: Explainable vision transformer based covid-19 screening using radiography. 2021.
- [231] Sangjoon Park, Gwanghyun Kim, Yujin Oh, Joon Beom Seo, Sang Min Lee, Jin Hwan Kim, Sungjun Moon, Jae-Kwang Lim, and Jong Chul Ye. Vision transformer for covid-19 cxr diagnosis using chest x-ray feature corpus. *arXiv preprint arXiv:2103.07055*, 2021.
- [232] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020.
- [233] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019.
- [234] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 782–791, 2021.
- [235] Xiaohong Gao, Yu Qian, and Alice Gao. Covid-vit: Classification of covid-19 from ct chest images based on vision transformer models. *arXiv preprint arXiv:2107.01682*, 2021.
- [236] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [237] Hayden Gunraj, Ali Sabri, David Koff, and Alexander Wong. Covid-net ct-2: Enhanced deep neural networks for detection of covid-19 from chest ct images through bigger, more diverse learning. *arXiv preprint arXiv:2101.07433*, 2021.
- [238] Debaditya Shome, T Kar, Sachi Nandan Mohanty, Prayag Tiwari, Khan Muhammad, Abdullah AlTameem, Yazhou Zhang, and Abdul Khader Jilani Saudagar. Covid-transformer: Interpretable covid-19 detection using vision transformer for healthcare. *International Journal of Environmental Research and Public Health*, 18(21):11086, 2021.
- [239] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, and Yongbing Zhang. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *arXiv preprint arXiv:2106.00908*, 2021.
- [240] Yin Dai, Yifan Gao, and Fayu Liu. Transmed: Transformers advance multi-modal medical image classification. *Diagnostics*, 11(8):1384, 2021.
- [241] Mengkang Lu, Yongsheng Pan, Dong Nie, Feihong Liu, Feng Shi, Yong Xia, and Dinggang Shen. Smile: Sparse-attention based multiple instance contrastive learning for glioma sub-type classification using pathological images. In *MICCAI Workshop on Computational Pathology*, pages 159–169. PMLR, 2021.

- [242] Sandy Napel and Sylvia K Plevritis. Nslc radiogenomics: initial stanford study of 26 cases. *Cancer Imaging Arch*, 2014.
- [243] Behnaz Gheflati and Hassan Rivaz. Vision transformers for classification of breast ultrasound images. *arXiv preprint arXiv:2110.14731*, 2021.
- [244] Walid Al-Dhabyani, Mohammed Gomaa, Hussien Khaled, and Aly Fahmy. Dataset of breast ultrasound images. *Data in brief*, 28:104863, 2020.
- [245] Moi Hoon Yap, Gerard Pons, Joan Martí, Sergi Ganau, Melcior Sentís, Reyer Zwiggelaar, Adrian K Davison, and Robert Martí. Automated breast ultrasound lesions detection using convolutional neural networks. *IEEE journal of biomedical and health informatics*, 22(4):1218–1226, 2017.
- [246] Anwar Khan and Boreom Lee. Gene transformer: Transformers for the gene expression-based classification of lung cancer subtypes. *arXiv preprint arXiv:2108.11833*, 2021.
- [247] Haoyuan Chen, Chen Li, Xiaoyan Li, Ge Wang, Weiming Hu, Yixin Li, Wanli Liu, Changhao Sun, Yudong Yao, Yueyang Teng, et al. Gashis-transformer: A multi-scale visual transformer approach for gastric histopathology image classification. *arXiv preprint arXiv:2104.14528*, 2021.
- [248] Zhencun Jiang, Zhengxin Dong, Lingyang Wang, and Wenping Jiang. Method for diagnosis of acute lymphoblastic leukemia based on vit-cnn ensemble model. *Computational Intelligence and Neuroscience*, 2021, 2021.
- [249] Glenn Fung, Murat Dundar, Balaji Krishnapuram, and R Bharat Rao. Multiple instance learning for computer aided diagnosis. *Advances in neural information processing systems*, 19:425, 2007.
- [250] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermesen, Quirine F Manson, Maschenka Balkenhol, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama*, 318(22):2199–2210, 2017.
- [251] The cancer genome atlas program. <https://www.cancer.gov/about-nci/organization/ccg/research/structural-genomics/tcga>. Accessed: 2022-01-20.
- [252] Yi Zheng, Rushin Gindra, Margrit Betke, Jennifer Beane, and Vijaya B Kolachalama. A deep learning based graph-transformer for whole slide image classification. *medRxiv*, 2021.
- [253] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [254] Shaimaa Bakr, Olivier Gevaert, Sebastian Echegaray, Kelsey Ayers, Mu Zhou, Majid Shafiq, Hong Zheng, Jalen Anthony Benson, Weiruo Zhang, Ann NC Leung, et al. A radiogenomic dataset of non-small cell lung cancer. *Scientific data*, 5(1):1–9, 2018.
- [255] Aptos 2019 blindness detection: Detect diabetic retinopathy to stop blindness before it's too late. <https://www.kaggle.com/c/aptos2019-blindness-detection>. Accessed: 2022-01-20.
- [256] Skin lesion analysis towards melanoma detection, 2019. <https://challenge2018.isic-archive.com/>. Accessed: 2022-01-20.
- [257] Rebecca Sawyer Lee, Francisco Gimenez, Assaf Hoogi, Kanae Kawai Miyake, Mia Gorovoy, and Daniel L Rubin. A curated mammography data set for use in computer-aided detection and diagnosis research. *Scientific data*, 4(1):1–9, 2017.
- [258] Shuang Yu, Kai Ma, Qi Bi, Cheng Bian, Munan Ning, Nanjun He, Yuexiang Li, Hanruo Liu, and Yefeng Zheng. Mil-vt: Multiple instance learning enhanced vision transformer for fundus image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 45–54. Springer, 2021.
- [259] Gwenolé Quéllec, Mathieu Lamard, Pierre-Henri Conze, Pascale Massin, and Béatrice Cochener. Automatic detection of rare pathologies in fundus photographs using few-shot learning. *Medical image analysis*, 61:101660, 2020.
- [260] Rui Sun, Yihao Li, Tianzhu Zhang, Zhendong Mao, Feng Wu, and Yongdong Zhang. Lesion-aware transformers for diabetic retinopathy grading. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10938–10947, 2021.
- [261] Etienne Decencière, Xiwei Zhang, Guy Cazuguel, Bruno Lay, Béatrice Cochener, Caroline Trone, Philippe Gain, Richard Ordonez, Pascale Massin, Ali Erginay, et al. Feedback on a publicly distributed image database: the messidor database. *Image Analysis & Stereology*, 33(3):231–234, 2014.
- [262] Jorge Cuadros and George Bresnick. Eyepacs: an adaptable telemedicine system for diabetic retinopathy screening. *Journal of diabetes science and technology*, 3(3):509–516, 2009.
- [263] Siim-acr pneumothorax segmentation. <https://www.kaggle.com/c/siim-acr-pneumothorax-segmentation>. Accessed: 2022-01-20.
- [264] Rsna pneumonia detection challenge (2018). <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge>. Accessed: 2022-01-20.
- [265] Linda Wang, Zhong Qiu Lin, and Alexander Wong. Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Scientific Reports*, 10(1):1–12, 2020.
- [266] Ara Abigail E Ambita, Eujene Nikka V Boquio, and Prospero C Naval. Covit-gan: Vision transformer for covid-19 detection in ct scan images with self-attention gan for data augmentation. In *International Conference on Artificial Neural Networks*, pages 587–598. Springer, 2021.
- [267] Plamen Angelov and Eduardo Almeida Soares. Sars-cov-2 ct-scan dataset: A large dataset of real patients ct scans for sars-cov-2 identification. *MedRxiv*, 2020.
- [268] Walid El-Shafaï and Fathi Abd El-Samie. Extensive covid-19 x-ray and ct chest images dataset. *mendeley data*, v3, 2020.
- [269] Unais Sait, KG Lal, S Prajapati, R Bhaumik, T Kumar, S Sanjana, and K Bhalla. Curated dataset for covid-19 posterior-anterior chest radiography images (x-rays). *Mendeley Data*, 1, 2020.
- [270] Xiao Qi, Lloyd G Brown, David J Foran, John Noshier, and Ilker Hacihaliloglu. Chest x-ray image phase features for improved diagnosis of covid-19 using convolutional neural network. *International journal of computer assisted radiology and surgery*, 16(2):197–206, 2021.
- [271] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *arXiv preprint arXiv:2104.14294*, 2021.
- [272] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific data*, 5(1):1–9, 2018.
- [273] Honggang Yang, Jiejie Chen, and Mengfei Xu. Fundus disease image classification based on improved transformer. In *2021 International Conference on Neuromorphic Computing (ICNC)*, pages 207–214. IEEE, 2021.
- [274] Ophthalmic image analysis dataset. <https://github.com/nkicsl/OIA>. Accessed: 2022-01-20.
- [275] Jianfang Wu, Ruo Hu, Zhenghong Xiao, Jiaxu Chen, and Jingwei Liu. Vision transformer-based recognition of diabetic retinopathy grade. *Medical Physics*, 2021.
- [276] Nouar AIDahoul, Hezerul Abdul Karim, Myles Joshua Toledo Tan, Mhd Adel Momo, and Jamie Ledesma Fermin. Encoding retina image to words using ensemble of vision transformers for diabetic retinopathy grading. *F1000Research*, 10:948, 2021.
- [277] Yutong Xie, Jianpeng Zhang, Yong Xia, and Qi Wu. Unified 2d and 3d pre-training for medical image classification and segmentation. *arXiv preprint arXiv:2112.09356*, 2021.
- [278] Fangzhou Liao, Ming Liang, Zhe Li, Xiaolin Hu, and Sen Song. Evaluate the malignancy of pulmonary nodules using the 3-d deep leaky noisy-or network. *IEEE transactions on neural networks and learning systems*, 30(11):3484–3495, 2019.
- [279] Nilay Ganatra. A comprehensive study of applying object detection methods for medical image analysis. In *2021 8th International Conference on Computing for Sustainable Global Development (INDIACom)*, pages 821–826. IEEE, 2021.
- [280] Shijie Liu, Hongyu Zhou, Xiaozhou Shi, and Junwen Pan. Transformer for polyp detection. *arXiv preprint arXiv:2111.07918*, 2021.
- [281] Tejas Sudharshan Mathai, Sungwon Lee, Daniel C Elton, Thomas C Shen, Yifan Peng, Zhiyong Lu, and Ronald M Summers. Lymph node detection in t2 mri with transformers. *arXiv preprint arXiv:2111.04885*, 2021.
- [282] Muhammad Maaz, Hanoona Rasheed, Salman Khan, Fahad Shahbaz Khan, Rao Muhammad Anwer, and Ming-Hsuan Yang. Multi-modal transformers excel at class-agnostic object detection. *arXiv preprint arXiv:2111.11430*, 2021.
- [283] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021.

- [284] Muralikrishna Gopal, Shaad E Abdullah, James J Grady, and James S Goodwin. Screening for lung cancer with low-dose computed tomography: a systematic review and meta-analysis of the baseline findings of randomized controlled trials. *Journal of thoracic oncology*, 5(8):1233–1239, 2010.
- [285] Alexandre Sadate, Bob V Océan, Jean-Paul Beregi, Aymeric Hamard, Takieddine Addala, Hélène de Forges, Pascale Fabbro-Peray, and Julien Frandon. Systematic review and meta-analysis on the impact of lung cancer screening by low-dose computed tomography. *European Journal of Cancer*, 134:107–114, 2020.
- [286] Cynthia H McCollough, Adam C Bartley, Rickey E Carter, Baiyu Chen, Tammy A Drees, Phillip Edwards, David R Holmes III, Alice E Huang, Farhana Khan, Shuai Leng, et al. Low-dose ct for the detection and classification of metastatic liver lesions: results of the 2016 low dose ct grand challenge. *Medical physics*, 44(10):e339–e352, 2017.
- [287] Jeffrey A Fessler. Model-based image reconstruction for mri. *IEEE signal processing magazine*, 27(4):81–89, 2010.
- [288] Dominik Narnhofer, Kerstin Hammernik, Florian Knoll, and Thomas Pock. Inverse gans for accelerated mri reconstruction. In *Wavelets and Sparsity XVIII*, volume 11138, page 111381A. International Society for Optics and Photonics, 2019.
- [289] Dayang Wang, Zhan Wu, and Hengyong Yu. Ted-net: Convolution-free t2t vision transformer-based encoder-decoder dilation network for low-dose ct denoising. *arXiv preprint arXiv:2106.04650*, 2021.
- [290] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zihang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *arXiv preprint arXiv:2101.11986*, 2021.
- [291] Achleshwar Luthra, Harsh Sulakhe, Tanish Mittal, Abhishek Iyer, and Santosh Yadav. Eformer: Edge enhancement based transformer for medical image denoising. *arXiv preprint arXiv:2109.08044*, 2021.
- [292] Tengfei Liang, Yi Jin, Yidong Li, and Tao Wang. Edcnn: Edge enhancement-based densely connected network with compound loss for low-dose ct denoising. In *2020 15th IEEE International Conference on Signal Processing (ICSP)*, volume 1, pages 193–198. IEEE, 2020.
- [293] FG Irwin et al. An isotropic 3x3 image gradient operator. Presentation at Stanford AI Project, 2014(02), 1968.
- [294] Yanmei Luo, Yan Wang, Chen Zu, Bo Zhan, Xi Wu, Jiliu Zhou, Dinggang Shen, and Luping Zhou. 3d transformer-gan for high-quality pet reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 276–285. Springer, 2021.
- [295] Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, Sungchul Lee, and Jin Keun Seo. Deep learning for undersampled mri reconstruction. *Physics in Medicine & Biology*, 63(13):135007, 2018.
- [296] Chun-Mei Feng, Yunlu Yan, Geng Chen, Huazhu Fu, Yong Xu, and Ling Shao. Accelerated multi-modal mr imaging with transformers. *arXiv preprint arXiv:2106.14248*, 2021.
- [297] Chun-Mei Feng, Yunlu Yan, Huazhu Fu, Li Chen, and Yong Xu. Task transformer network for joint mri reconstruction and super-resolution. *arXiv preprint arXiv:2106.06742*, 2021.
- [298] Dwarikanath Mahapatra and Zongyuan Ge. Mr image super resolution by combining feature disentanglement cnns and vision transformers. 2021.
- [299] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. *arXiv preprint arXiv:2007.00653*, 2020.
- [300] Yilmaz Korkmaz, Mahmut Yurt, Salman Ul Hassan Dar, Muzaffer Özbey, and Tolga Cukur. Deep mri reconstruction with generative vision transformers. In *International Workshop on Machine Learning for Medical Image Reconstruction*, pages 54–64. Springer, 2021.
- [301] Kang Lin and Reinhard Heckel. Vision transformers enable fast and robust accelerated mri. 2021.
- [302] Hamid R Sheikh and Alan C Bovik. Image information and visual quality. *IEEE Transactions on image processing*, 15(2):430–444, 2006.
- [303] Information extraction from images (ixi). <http://brain-development.org/ixi-dataset/>. Accessed: 2022-01-20.
- [304] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv preprint arXiv:1811.08839*, 2018.
- [305] Ce Wang, Kun Shang, Haimiao Zhang, Qian Li, Yuan Hui, and S Kevin Zhou. Dudotrans: Dual-domain transformer provides more attention for sinogram restoration in sparse-view ct reconstruction. *arXiv preprint arXiv:2111.10790*, 2021.
- [306] Jiayi Pan, Weiwen Wu, Zhifan Gao, and Heye Zhang. Multi-domain integrative swin transformer network for sparse-view tomographic reconstruction. Available at SSRN 3991087.
- [307] Yonghao Long, Zhaoshuo Li, Chi Hang Yee, Chi Fai Ng, Russell H Taylor, Mathias Unberath, and Qi Dou. E-dssr: Efficient dynamic surgical scene reconstruction with transformer-based stereoscopic depth perception. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 415–425. Springer, 2021.
- [308] Menglong Ye, Edward Johns, Ankur Handa, Lin Zhang, Philip Pratt, and Guang-Zhong Yang. Self-supervised siamese learning on stereo image pairs for depth estimation in robotic surgery. *arXiv preprint arXiv:1705.08260*, 2017.
- [309] Alper Güngör, Baris Askin, Damla Alptekin Soydan, Emine Ulku Saritas, Can Barış Top, and Tolga Çukur. Transms: Transformers for super-resolution calibration in magnetic particle imaging. *arXiv preprint arXiv:2111.02163*, 2021.
- [310] Tobias Knopp, Patryk Szwargulski, Florian Griesse, and Matthias Gräser. Openmpidata: An initiative for freely accessible magnetic particle imaging data. *Data in brief*, 28:104971, 2020.
- [311] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated mri reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 64–73. Springer, 2020.
- [312] Yoseob Han and Jong Chul Ye. Framing u-net via deep convolutional framelets: Application to sparse-view ct. *IEEE transactions on medical imaging*, 37(6):1418–1429, 2018.
- [313] Hiroyuki Kudo, Taizo Suzuki, and Essam A Rashed. Image reconstruction for sparse-view ct and interior ct—introduction to compressed sensing and differentiated backprojection. *Quantitative imaging in medicine and surgery*, 3(3):147, 2013.
- [314] Chen Liu, Mathieu Salzmann, Tao Lin, Ryota Tomioka, and Sabine Süsstrunk. On the loss landscape of adversarial training: Identifying challenges and how to overcome them. *arXiv preprint arXiv:2006.08403*, 2020.
- [315] Xuzhe Zhang, Xinzi He, Jia Guo, Nabil Ettehad, Natalie Aw, David Semanek, Jonathan Posner, Andrew Laine, and Yun Wang. Ptnet: A high-resolution infant mri synthesizer based on transformer. *arXiv preprint arXiv:2105.13993*, 2021.
- [316] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [317] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.
- [318] Antonios Makropoulos, Emma C Robinson, Andreas Schuh, Robert Wright, Sean Fitzgibbon, Jelena Bozek, Serena J Counsell, Johannes Steinweg, Katy Vecchiato, Jonathan Passerat-Palmbach, et al. The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. *Neuroimage*, 173:88–112, 2018.
- [319] Sharif Amit Kamran, Khondker Fariha Hossain, Alireza Tavakkoli, Stewart Lee Zuckerbrod, Kenton M Sanders, and Salah A Baker. Vtgan: Semi-supervised retinal image synthesis and disease prediction using vision transformers. *arXiv preprint arXiv:2104.06757*, 2021.
- [320] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [321] Mikolaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.

- [322] Shirin Hajeb Mohammad Alipour, Hossein Rabbani, and Mohammad Reza Akhlaghi. Diabetic retinopathy grading by digital curvelet transform. *Computational and mathematical methods in medicine*, 2012, 2012.
- [323] Nicolae-Catalin Ristea, Andreea-Iuliana Miron, Olivian Savencu, Mariana-Iuliana Georgescu, Nicolae Verga, Fahad Shahbaz Khan, and Radu Tudor Ionescu. Cytran: Cycle-consistent transformers for non-contrast to contrast ct translation. *arXiv preprint arXiv:2110.06400*, 2021.
- [324] Jelmer M Wolterink, Anna M Dinkla, Mark HF Savenije, Peter R Seevinck, Cornelis AT van den Berg, and Ivana Išgum. Deep mr to ct synthesis using unpaired data. In *International workshop on simulation and synthesis in medical imaging*, pages 14–23. Springer, 2017.
- [325] Onat Dalmaz, Mahmut Yurt, and Tolga Çukur. Resvit: Residual vision transformers for multi-modal medical image synthesis. *arXiv preprint arXiv:2106.16031*, 2021.
- [326] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- [327] Tufve Nyholm, Stina Svensson, Sebastian Andersson, Joakim Jonsson, Maja Sohlén, Christian Gustafsson, Elisabeth Kjellén, Karin Söderström, Per Albertsson, Lennart Blomqvist, et al. Mr and ct data with multiobserver delineations of organs in the pelvic area—part of the gold atlas project. *Medical physics*, 45(3):1295–1300, 2018.
- [328] Zhengwei Wang, Qi She, and Tomas E Ward. Generative adversarial networks in computer vision: A survey and taxonomy. *ACM Computing Surveys (CSUR)*, 54(2):1–38, 2021.
- [329] Fakhre Alam, Sami Ur Rahman, Sehat Ullah, and Kamal Gulati. Medical image registration in image guided surgery: Issues, challenges and research opportunities. *Biocybernetics and Biomedical Engineering*, 38(1):71–89, 2018.
- [330] Junyu Chen, Yufan He, Eric C Frey, Ye Li, and Yong Du. Vit-v-net: Vision transformer for unsupervised volumetric medical image registration. *arXiv preprint arXiv:2104.06468*, 2021.
- [331] F Milletari, N Navab, and SA Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. *June 2016*, 2018.
- [332] Junyu Chen, Yong Du, Yufan He, William P Segars, Ye Li, and Eric C Frey. Transmorph: Transformer for unsupervised medical image registration. *arXiv preprint arXiv:2111.10480*, 2021.
- [333] Yungeng Zhang, Yuru Pei, and Hongbin Zha. Learning dual transformer network for diffeomorphic registration. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 129–138. Springer, 2021.
- [334] Daniel S Marcus, Tracy H Wang, Jamie Parker, John G Csernansky, John C Morris, and Randy L Buckner. Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of cognitive neuroscience*, 19(9):1498–1507, 2007.
- [335] John Pavlopoulos, Vasiliki Kougia, Ion Androutsopoulos, and Dimitris Papamichail. Diagnostic captioning: a survey. *arXiv preprint arXiv:2101.07299*, 2021.
- [336] Maram Mahmoud A Monshi, Josiah Poon, and Vera Chung. Deep learning in generating radiology reports: A survey. *Artificial Intelligence in Medicine*, 106:101878, 2020.
- [337] Vasiliki Kougia, John Pavlopoulos, and Ion Androutsopoulos. A survey on biomedical image captioning. *arXiv preprint arXiv:1905.13302*, 2019.
- [338] Pablo Messina, Pablo Pino, Denis Parra, Alvaro Soto, Cecilia Besa, Sergio Uribe, Cristian Tejos, Claudia Prieto, Daniel Capurro, et al. A survey on deep learning and explainability for automatic image-based medical report generation. *arXiv preprint arXiv:2010.10563*, 2020.
- [339] Steven J Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. Self-critical sequence training for image captioning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7008–7024, 2017.
- [340] Yuxuan Xiong, Bo Du, and Pingkun Yan. Reinforced transformer for medical image captioning. In *International Workshop on Machine Learning in Medical Imaging*, pages 673–680. Springer, 2019.
- [341] Dina Demner-Fushman, Marc D Kohli, Marc B Rosenman, Sonya E Shooshan, Laritza Rodriguez, Sameer Antani, George R Thoma, and Clement J McDonald. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association*, 2016.
- [342] Yasuhide Miura, Yuhao Zhang, Emily Bao Tsai, Curtis P Langlotz, and Dan Jurafsky. Improving factual completeness and consistency of image-to-text radiology report generation. *arXiv preprint arXiv:2010.10042*, 2020.
- [343] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
- [344] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015.
- [345] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*, 2019.
- [346] Jinglu Zhang, Yinyu Nie, Jian Chang, and Jian Jun Zhang. Surgical instruction generation with transformers. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 290–299. Springer, 2021.
- [347] Edgar Rojas-Muñoz, Kyle Couperus, and Juan Wachs. Daisi: Database for ai surgical instruction. *arXiv preprint arXiv:2004.02809*, 2020.
- [348] Preethi Srinivasan, Daksh Thapar, Arnav Bhavsar, and Aditya Nigam. Hierarchical x-ray report generation via pathology tags and multi head attention. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [349] Fenglin Liu, Xian Wu, Shen Ge, Wei Fan, and Yuexian Zou. Exploring and distilling posterior and prior knowledge for radiology report generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13753–13762, 2021.
- [350] Di You, Fenglin Liu, Shen Ge, Xiaoxia Xie, Jing Zhang, and Xian Wu. Aligntransformer: Hierarchical alignment of visual regions and disease tags for medical report generation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 72–82. Springer, 2021.
- [351] Marcella Cornia, Matteo Stefanini, Lorenzo Baraldi, and Rita Cucchiara. Meshed-memory transformer for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10578–10587, 2020.
- [352] Alistair EW Johnson, Tom J Pollard, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Yifan Peng, Zhiyong Lu, Roger G Mark, Seth J Berkowitz, and Steven Horng. Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042*, 2019.
- [353] Zhihong Chen, Yaling Shen, Yan Song, and Xiang Wan. Cross-modal memory networks for radiology report generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5904–5914, 2021.
- [354] Mingjie Li, Wenjia Cai, Rui Liu, Yuetian Weng, Xiaoyun Zhao, Cong Fang, Xin Chen, Zhong Liu, Caineng Pan, Mengke Li, et al. Ff-ir: Towards an explainable and reliable medical report generation benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [355] Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. Padchest: A large chest x-ray image dataset with multi-label annotated reports. *Medical image analysis*, 66:101797, 2020.
- [356] Christy Y Li, Xiaodan Liang, Zhiting Hu, and Eric P Xing. Knowledge-driven encode, retrieve, paraphrase for medical image report generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6666–6673, 2019.
- [357] RVJPH Kälviäinen and H Uusitalo. Diaretdb1 diabetic retinopathy database and evaluation protocol. In *Medical Image Understanding and Analysis*, volume 2007, page 61. Citeseer, 2007.
- [358] Mingjie Li, Fuyu Wang, Xiaojun Chang, and Xiaodan Liang. Auxiliary signal-guided knowledge encoder-decoder for medical report generation. *arXiv preprint arXiv:2006.03744*, 2020.
- [359] Jia-Hong Huang, C-H Huck Yang, Fangyu Liu, Meng Tian, Yi-Chieh Liu, Ting-Wei Wu, I Lin, Kang Wang, Hiromasa Morikawa,

- Hernghua Chang, et al. Deepopht: medical report generation for retinal images via deep models and visual explanation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2442–2452, 2021.
- [360] AD Hoover, Valentina Kouznetsova, and Michael Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Transactions on Medical imaging*, 19(3):203–210, 2000.
- [361] An Yan, Zexue He, Xing Lu, Jiang Du, Eric Chang, Amilcare Gentili, Julian McAuley, and Chun-Nan Hsu. Weakly supervised contrastive learning for chest x-ray report generation. *arXiv preprint arXiv:2109.12242*, 2021.
- [362] Ayoub Benali Amjoud and Mustapha Amrouch. Automatic generation of chest x-ray reports using a transformer-based deep learning model. In *2021 Fifth International Conference On Intelligent Computing in Data Sciences (ICDS)*, pages 1–5. IEEE, 2021.
- [363] Benjamin Hou, Georgios Kaissis, Ronald M Summers, and Bernhard Kainz. Ratchet: Medical transformer for chest x-ray diagnosis and reporting. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 293–303. Springer, 2021.
- [364] Zhihong Chen, Yan Song, Tsung-Hui Chang, and Xiang Wan. Generating radiology reports via memory-driven transformer. *arXiv preprint arXiv:2010.16056*, 2020.
- [365] Farhad Nooralahzadeh, Nicolas Perez Gonzalez, Thomas Frauenfelder, Koji Fujimoto, and Michael Krauthammer. Progressive transformer-based generation of radiology reports. *arXiv preprint arXiv:2102.09777*, 2021.
- [366] Yixin Wang, Zihao Lin, Jiang Tian, Zhongchao Shi, Yang Zhang, Jianping Fan, and Zhiqiang He. Confidence-guided radiology report generation. *arXiv preprint arXiv:2106.10887*, 2021.
- [367] Hoang TN Nguyen, Dong Nie, Taivanbat Badamdorj, Yujie Liu, Yingying Zhu, Jason Truong, and Li Cheng. Automated generation of accurate & fluent medical x-ray reports. *arXiv preprint arXiv:2108.12126*, 2021.
- [368] Fenglin Liu, Chenyu You, Xian Wu, Shen Ge, Xu Sun, et al. Auto-encoding knowledge graph for unsupervised medical report generation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [369] Justin Lovelace and Bobak Mortazavi. Learning to generate clinically coherent chest x-ray reports. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 1235–1243, 2020.
- [370] Michael Denkowski and Alon Lavie. Meteor 1.3: Automatic metric for reliable optimization and evaluation of machine translation systems. In *Proceedings of the sixth workshop on statistical machine translation*, pages 85–91, 2011.
- [371] Lin CY ROUGE. A package for automatic evaluation of summaries. In *Proceedings of Workshop on Text Summarization of ACL, Spain*, 2004.
- [372] Omar Alfarghaly, Rana Khaled, Abeer Elkorany, Maha Helal, and Aly Fahmy. Automated radiology report generation using conditioned transformers. *Informatics in Medicine Unlocked*, 24:100557, 2021.
- [373] Zachary M Ziegler, Luke Melas-Kyriazi, Sebastian Gehrmann, and Alexander M Rush. Encoder-agnostic adaptation for conditional language generation. *arXiv preprint arXiv:1908.06938*, 2019.
- [374] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [375] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48, 2009.
- [376] Esha Pahwa, Dwij Mehta, Sanjeet Kapadia, Devansh Jain, and Achleshwar Luthra. Medskip: Medical report generation using skip connections and integrated attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3409–3415, 2021.
- [377] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5693–5703, 2019.
- [378] Hyeryun Park, Kyungmo Kim, Seongkeun Park, and Jinwook Choi. Medical image captioning model to convey more details: Methodological comparison of feature difference generation. *IEEE Access*, 9:150560–150568, 2021.
- [379] Obioma Pelka, Sven Koitka, Johannes Rückert, Felix Nensa, and Christoph M Friedrich. Radiology objects in context (roco): a multimodal image dataset. In *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*, pages 180–189. Springer, 2018.
- [380] Baoyu Jing, Pengtao Xie, and Eric Xing. On the automatic generation of medical imaging reports. *arXiv preprint arXiv:1711.08195*, 2017.
- [381] Alba Garcia Seco De Herrera, Carstern Eickhof, Vincent Andrearczyk, and Henning Müller. Overview of the imageclef 2018 caption prediction tasks. *CEUR Workshop Proceedings*, 2018.
- [382] Richard J Chen, Ming Y Lu, Wei-Hung Weng, Tiffany Y Chen, Drew FK Williamson, Trevor Manz, Maha Shady, and Faisal Mahmood. Multimodal co-attention transformer for survival prediction in gigapixel whole slide images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4025, 2021.
- [383] Elly Kipkogei, Arango Argoty GA, Ioannis Kagiampakis, Arijit Patra, and Etai Jacob. Explainable transformer-based neural network for the prediction of survival outcomes in non-small cell lung cancer (nsccl). 2021.
- [384] Robert M Samstein, Chung-Han Lee, Alexander N Shoushtari, Matthew D Hellmann, Ronglai Shen, Yelena Y Janjigian, David A Barron, Ahmet Zehir, Emmet J Jordan, Antonio Omuro, et al. Tumor mutational load predicts survival after immunotherapy across multiple cancer types. *Nature genetics*, 51(2):202–206, 2019.
- [385] Sedigheh Eslami, de Gerard Melo, and Christoph Meinel. Does clip benefit visual question answering in the medical domain as much as it does in the general domain? *arXiv preprint arXiv:2112.13906*, 2021.
- [386] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021.
- [387] Itzik Malkiel, Gony Rosenman, Lior Wolf, and Talma Hendler. Pre-training and fine-tuning transformers for fmri prediction tasks. *arXiv preprint arXiv:2112.05761*, 2021.
- [388] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent: A new approach to self-supervised learning. *arXiv preprint arXiv:2006.07733*, 2020.
- [389] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021.
- [390] 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. *arXiv preprint arXiv:2101.05224*, 2021.
- [391] Yen Nhi Truong Vu, Richard Wang, Niranjana Balachandrar, Can Liu, Andrew Y Ng, and Pranav Rajpurkar. Medaug: Contrastive learning leveraging patient metadata improves representations for chest x-ray interpretation. *arXiv preprint arXiv:2102.10663*, 2021.
- [392] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *arXiv preprint arXiv:2006.10029*, 2020.
- [393] Mauricio Reyes, Raphael Meier, Sérgio Pereira, Carlos A Silva, Fried-Michael Dahlweid, Hendrik von Tengg-Koblick, Ronald M Summers, and Roland Wiest. On the interpretability of artificial intelligence in radiology: challenges and opportunities. *Radiology: Artificial Intelligence*, 2(3):e190043, 2020.
- [394] Amitojdeep Singh, Sourya Sengupta, and Vasudevan Lakshminarayanan. Explainable deep learning models in medical image analysis. *Journal of Imaging*, 6(6):52, 2020.
- [395] Ke Yan, Youbao Tang, Adam P Harrison, Jinzheng Cai, Le Lu, and Jingjing Lu. Interpretable medical image classification with self-supervised anatomical embedding and prior knowledge. 2021.
- [396] Biraja Ghoshal and Allan Tucker. Estimating uncertainty and interpretability in deep learning for coronavirus (covid-19) detection. *arXiv preprint arXiv:2003.10769*, 2020.
- [397] Sneha Chaudhari, Varun Mithal, Gungor Polatkan, and Rohan Ramanath. An attentive survey of attention models.

- ACM Transactions on Intelligent Systems and Technology (TIST), 12(5):1–32, 2021.
- [398] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, Alessandro Crimi, Russell Takeshi Shinohara, Christoph Berger, Sung Min Ha, Martin Rozycki, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*, 2018.
- [399] Open access series of imaging studies. <https://www.oasis-brains.org/>. Accessed: 2022-01-20.
- [400] Lung nodule analysis, 2016. <https://luna16.grand-challenge.org/Data/>. Accessed: 2022-01-20.
- [401] Daniel S Kermany, Michael Goldbaum, Wenjia Cai, Carolina CS Valentim, Huiying Liang, Sally L Baxter, Alex McKeown, Ge Yang, Xiaokang Wu, Fangbing Yan, et al. Identifying medical diagnoses and treatable diseases by image-based deep learning. *Cell*, 172(5):1122–1131, 2018.
- [402] National lung screening trial - the cancer data assess system. <https://cdas.cancer.gov/nlst/>. Accessed: 2022-01-20.
- [403] Diabetic retinopathy challenge. <https://www.kaggle.com/c/diabetic-retinopathy-detection>. Accessed: 2022-01-20.
- [404] Xingjun Ma, Yuhao Niu, Lin Gu, Yisen Wang, Yitian Zhao, James Bailey, and Feng Lu. Understanding adversarial attacks on deep learning based medical image analysis systems. *Pattern Recognition*, 110:107332, 2021.
- [405] Konstantinos Papangelou, Konstantinos Sechidis, James Weatherall, and Gavin Brown. Toward an understanding of adversarial examples in clinical trials. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 35–51. Springer, 2018.
- [406] Samuel G Finlayson, John D Bowers, Joichi Ito, Jonathan L Zittrain, Andrew L Beam, and Isaac S Kohane. Adversarial attacks on medical machine learning. *Science*, 363(6433):1287–1289, 2019.
- [407] Philipp Benz, Soomin Ham, Chaoning Zhang, Adil Karjauv, and In So Kweon. Adversarial robustness comparison of vision transformer and mlp-mixer to cnns. *arXiv preprint arXiv:2110.02797*, 2021.
- [408] Srinadh Bhojanapalli, Ayan Chakrabarti, Daniel Glasner, Daliang Li, Thomas Unterthiner, and Andreas Veit. Understanding robustness of transformers for image classification. *arXiv preprint arXiv:2103.14586*, 2021.
- [409] Zhipeng Wei, Jingjing Chen, Micah Goldblum, Zuxuan Wu, Tom Goldstein, and Yu-Gang Jiang. Towards transferable adversarial attacks on vision transformers. *arXiv preprint arXiv:2109.04176*, 2021.
- [410] Rulin Shao, Zhouxing Shi, Jinfeng Yi, Pin-Yu Chen, and Cho-Jui Hsieh. On the adversarial robustness of visual transformers. *arXiv preprint arXiv:2103.15670*, 2021.
- [411] Sayak Paul and Pin-Yu Chen. Vision transformers are robust learners. *arXiv preprint arXiv:2105.07581*, 2021.
- [412] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Fahad Shahbaz Khan, and Fatih Porikli. On improving adversarial transferability of vision transformers. *arXiv preprint arXiv:2106.04169*, 2021.
- [413] Xiaofeng Mao, Gege Qi, Yuefeng Chen, Xiaodan Li, Shaokai Ye, Yuan He, and Hui Xue. Rethinking the design principles of robust vision transformer. *arXiv preprint arXiv:2105.07926*, 2021.
- [414] Kaleel Mahmood, Rigel Mahmood, and Marten Van Dijk. On the robustness of vision transformers to adversarial examples. *arXiv preprint arXiv:2104.02610*, 2021.
- [415] Ameya Joshi, Gauri Jagatap, and Chinmay Hegde. Adversarial token attacks on vision transformers. *arXiv preprint arXiv:2110.04337*, 2021.
- [416] Ahmed Aldahdooh, Wassim Hamidouche, and Olivier Deforges. Reveal of vision transformers robustness against adversarial attacks. *arXiv preprint arXiv:2106.03734*, 2021.
- [417] Guy Katz, Clark Barrett, David L Dill, Kyle Julian, and Mykel J Kochenderfer. Towards proving the adversarial robustness of deep neural networks. *arXiv preprint arXiv:1709.02802*, 2017.
- [418] K Bonawitz, H Eichner, W Grieskamp, et al. Tensorflow federated: Machine learning on decentralized data. 2020. <https://www.tensorflow.org/federated>.
- [419] Brian Knott, Shobha Venkataraman, Awni Hannun, Shubho Sengupta, Mark Ibrahim, and Laurens van der Maaten. Crypten: Secure multi-party computation meets machine learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [420] To lower the barrier to entry to privacy preserving technology. <https://github.com/OpenMined>. Accessed: 2022-01-20.
- [421] Ashkan Yousefpour, Igor Shilov, Alexandre Sablayrolles, Davide Testuggine, Karthik Prasad, Mani Malek, John Nguyen, Sayan Ghosh, Akash Bharadwaj, Jessica Zhao, et al. Opacus: User-friendly differential privacy library in pytorch. *arXiv preprint arXiv:2109.12298*, 2021.
- [422] Georgios Kaissis, Alexander Ziller, Jonathan Passerat-Palmbach, Théo Ryffel, Dmitrii Usynin, Andrew Trask, Ionésio Lima, Jason Mancuso, Friederike Jungmann, Marc-Matthias Steinborn, et al. End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence*, 3(6):473–484, 2021.
- [423] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. *arXiv preprint arXiv:2106.04560*, 2021.
- [424] Hamid Tabani, Ajay Balasubramaniam, Shabbir Marzban, Elahe Arani, and Bahram Zonooz. Improving the efficiency of transformers for resource-constrained devices. In *2021 24th Euromicro Conference on Digital System Design (DSD)*, pages 449–456. IEEE, 2021.
- [425] Bingbing Li, Santosh Pandey, Haowen Fang, Yanjun Lyv, Ji Li, Jieyang Chen, Mimi Xie, Lipeng Wan, Hang Liu, and Caiwen Ding. Ftrans: energy-efficient acceleration of transformers using fpga. In *Proceedings of the ACM/IEEE International Symposium on Low Power Electronics and Design*, pages 175–180, 2020.
- [426] Hanrui Wang, Zhonghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. Hat: Hardware-aware transformers for efficient natural language processing. *arXiv preprint arXiv:2005.14187*, 2020.
- [427] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- [428] Yangsibo Huang, Samyak Gupta, Zhao Song, Kai Li, and Sanjeev Arora. Evaluating gradient inversion attacks and defenses in federated learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [429] Kimin Lee, Kibok Lee, Honglak Lee, and Jinwoo Shin. A simple unified framework for detecting out-of-distribution samples and adversarial attacks. *Advances in neural information processing systems*, 31, 2018.
- [430] Stanislav Fort, Jie Ren, and Balaji Lakshminarayanan. Exploring the limits of out-of-distribution detection. *arXiv preprint arXiv:2106.03004*, 2021.
- [431] Jingkang Yang, Kaiyang Zhou, Yixuan Li, and Ziwei Liu. Generalized out-of-distribution detection: A survey. *arXiv preprint arXiv:2110.11334*, 2021.
- [432] Oliver Zhang, Jean-Benoit Delbrouck, and Daniel L Rubin. Out of distribution detection for medical images. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Perinatal Imaging, Placental and Preterm Image Analysis*, pages 102–111. Springer, 2021.
- [433] Chongzhi Zhang, Mingyuan Zhang, Shanghang Zhang, Daisheng Jin, Qiang Zhou, Zhongang Cai, Haiyu Zhao, Shuai Yi, Xianglong Liu, and Ziwei Liu. Delving deep into the generalization of vision transformers under distribution shifts. *arXiv preprint arXiv:2106.07617*, 2021.
- [434] Jasper Linmans, Jeroen van der Laak, and Geert Litjens. Efficient out-of-distribution detection in digital pathology using multi-head convolutional neural networks. In *MIDL*, pages 465–478, 2020.
- [435] Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- [436] Rajat Koner, Poulami Sinhamahapatra, Karsten Roscher, Stephan Günnemann, and Volker Tresp. Oodformer: Out-of-distribution detection transformer. *arXiv preprint arXiv:2107.08976*, 2021.
- [437] Haoxuan Qu, Hossein Rahmani, Li Xu, Bryan Williams, and Jun Liu. Recent advances of continual learning in computer vision: An overview. *arXiv preprint arXiv:2109.11369*, 2021.
- [438] 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.
- [439] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020.