# Advances in Medical Image Analysis with Vision Transformers: A Comprehensive Review

Reza Azad[1], Amirhossein Kazerouni[2], Moein Heidari[2], Ehsan Khodapanah Aghdam[3], Amirali Molaei[4], Yiwei Jia[1], Abin Jose[1], Rijo Roy[1], Dorit Merhof[†,5,6*]

[a]*Faculty of Electrical Engineering and Information Technology, RWTH Aachen University, Aachen, Germany*
[b]*School of Electrical Engineering, Iran University of Science and Technology, Tehran, Iran*
[c]*Department of Electrical Engineering, Shahid Beheshti University, Tehran, Iran*
[d]*School of Computer Engineering, Iran University of Science and Technology, Tehran, Iran*
[e]*Faculty of Informatics and Data Science, University of Regensburg, Regensburg, Germany*
[f]*Fraunhofer Institute for Digital Medicine MEVIS, Bremen, Germany*

## Abstract

The remarkable performance of the Transformer architecture in natural language processing has recently also triggered broad interest in Computer Vision. Among other merits, Transformers are witnessed as capable of learning long-range dependencies and spatial correlations, which is a clear advantage over convolutional neural networks (CNNs), which have been the de facto standard in Computer Vision problems so far. Thus, Transformers have become an integral part of modern medical image analysis. In this review, we provide an encyclopedic review of the applications of Transformers in medical imaging. Specifically, we present a systematic and thorough review of relevant recent Transformer literature for different medical image analysis tasks, including classification, segmentation, detection, registration, synthesis, and clinical report generation. For each of these applications, we investigate the novelty, strengths and weaknesses of the different proposed strategies and develop taxonomies highlighting key properties and contributions. Further, if applicable, we outline current benchmarks on different datasets. Finally, we summarize key challenges and discuss different future research directions. In addition, we have provided cited papers with their corresponding implementations in https://github.com/mindflow-institue/Awesome-Transformer.

*Keywords:* Transformers, Medical Image Analysis, Vision Transformers, Deep Neural Networks

## 1. Introduction

Convolutional neural networks (CNNs) have been an integral part of research in the field of medical image analysis for many years. By virtue of convolutional filters whose primary function is to learn and extract necessary features from medical images, a wealth of research has been dedicated to CNNs ranging from tumor detection and classification [1], detection of skin lesions [2, 3] to brain tumor segmentation [4], to name only a few. CNNs have also contributed significantly to the analysis of different imaging modalities in clinical medicine, including X-ray radiography, computed tomography (CT), magnetic resonance imaging (MRI), ultrasound (US), and digital pathology. Despite their outstanding performance, CNNs suffer from conceptual limitations and are innately unable to model explicit long-distance dependencies due to the limited receptive field of convolution kernels. Moreover, the convolutional operator suffers from the fact that at inference time, it applies fixed weights regardless of any changes to the visual input. To mitigate the aforementioned problems, there have been great research efforts to integrate attention mechanisms, which can be regarded as a dynamic weight adjustment process based on input features

to the seminal CNN-based structures to improve the non-local modeling capability [5, 6, 7].

To this end, Wang et al. [8] designed a non-local flexible building block, which can be plugged into multiple intermediate convolution layers. SENet [9] suggested a channel attention squeeze-and-excitation (SE) block, which collects global information in order to recalibrate each channel accordingly to create a more robust representation. Inspired by this line of research, there has been an overwhelming influx of models with attention variants proposed in the medical imaging field [10, 11, 12]. Although these attention mechanisms allow the modeling of full image contextual information, as the computational complexity of these approaches typically grows quadratically with respect to spatial size, they imply an intensive computational burden, thus making them inefficient in the case of medical images that are dense in pixel resolution [13]. Moreover, despite the fact that the combination of the attention mechanism with the convolutional operation leads to systematic performance gains, these models inevitably suffer from constraints in learning long-range interactions. Transformers [14] have demonstrated exemplary performance on a broad range of natural language processing (NLP) tasks, e.g., machine translation, text classification, and question answering. Inspired by the eminent success of Transformer architectures in the field of NLP, they have become a widely applied

---

*Corresponding author: Dorit Merhof, Tel.: +49 (941) 943-68509, E-mail: dorit.merhof@ur.de.
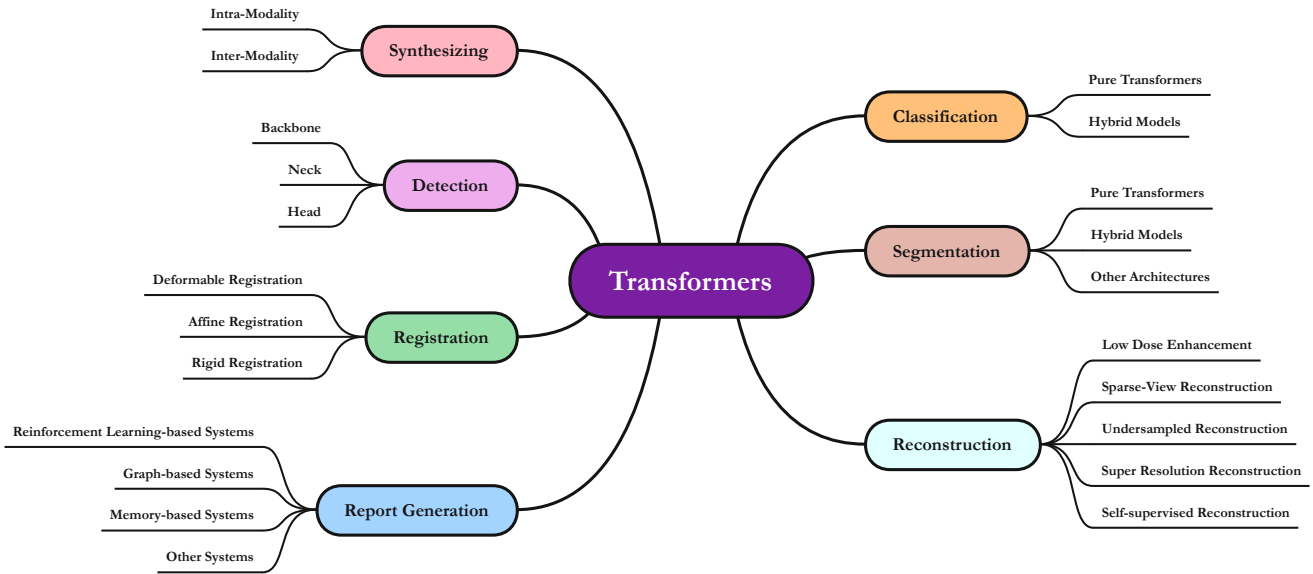
Figure 1: Overview of the applications covered in this review.

technique in modern Computer Vision (CV) models. Since the establishment of Vision-Transformers (ViTs) [15], Transformers proved to be valid alternatives to CNNs in diverse tasks ranging from image recognition [15], object detection [16], image segmentation [17] to video understanding [18] and image super-resolution [19]. As a central piece of the Transformer, the self-attention mechanism comes with the ability to model relationships between elements of a sequence, thereby learning long-range interactions [20]. Moreover, Transformers allow for large-scale pre-training for specific downstream tasks and applications and are capable of dealing with variable-length inputs. The immense interest in Transformers has also spurred research into medical imaging applications (see Figure 1). Being dominant in reputable top-tier medical imaging conferences and journals, it is extremely challenging for researchers and practitioners to keep up with the rate of innovation. The rapid adoption of Transformers in the medical imaging field necessitates a comprehensive summary and outlook, which is the main scope of this review. Specifically, this review provides a holistic overview of the Transformer models developed for medical imaging and image analysis applications. We provide a taxonomy of the network design, highlight the major strengths and deficiencies of the existing approaches, and introduce the current benchmarks in each task. We inspect several key technologies that arise from the various medical imaging applications, including medical image segmentation, medical image registration, medical image reconstruction, and medical image classification. So far, review papers related to Transformers do not concentrate on applications of Transformers in the medical imaging and image analysis domain [21]. The few literature reviews that do focus on the medical domain [22, 23], despite being very comprehensive, do not necessarily discuss the drawbacks and merits of each method. In our work, we explicitly cover this aspect and also provide a taxonomy that comprises the imaging modality, organ of interest, and type of

training procedure each paper has selected. More specifically, in Section 3 (Medical Image Classification), we comprehensively elaborate on the most promising networks along with their key ideas, limitations, the number of parameters, and the specific classification task they are addressing. In Section 4 (Medical Image Segmentation), we analyze network architectures in terms of their design choice and propose a detailed taxonomy to categorize each network to provide insight for the reader to understand the current limitations and progress in segmentation networks based on the Transformer architecture. In Section 5 (Medical Image Reconstruction), we take a different perspective to categorize networks based on their network structure and the imaging modality they are built upon. We categorize the synthesis methods in Section 6 based on their objective (intra-modality or inter-modality) and then provide detailed information regarding the network architecture, parameters, motivations, and highlights. In the sections related to detection (Section 7), registration (Section 8), and report generation (Section 9) we review in detail the state-of-the-art (SOTA) networks and provide detailed information regarding the network architectures, advantages, and drawbacks. Moreover, due to the swift development of the field, we believe that the community requires a more recent overview of the literature.

We hope this work will point out new research options and provide a guideline for researchers and initiate further interest in the vision community to leverage the potential of Transformer models in the medical domain. Our major contributions are as follows:

- We systematically and comprehensively review the applications of Transformers in the medical imaging domain and provide a comparison and analysis of SOTA approaches for each task. Specifically, more than 200 papers are covered in a hierarchical and structured manner.

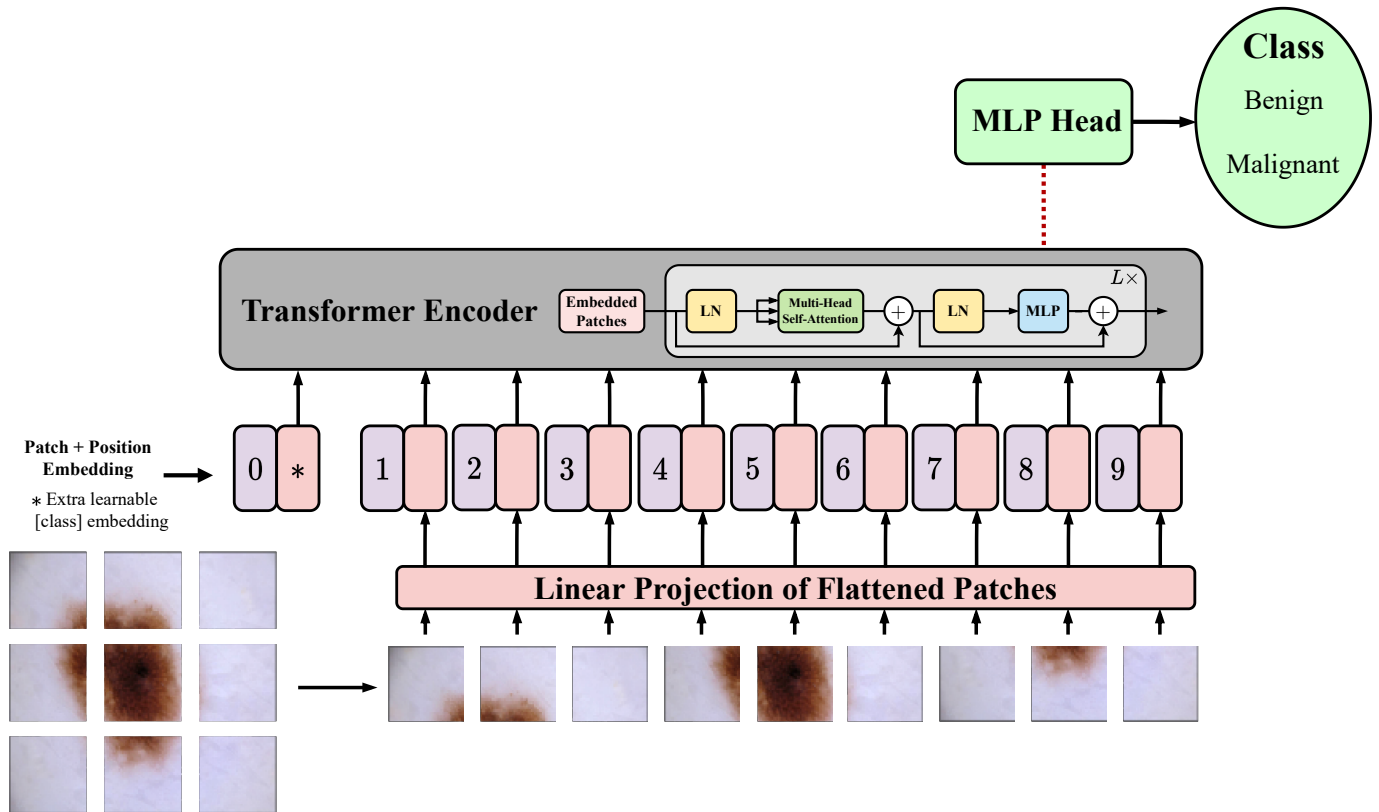- Our work provides a taxonomized (Figure 1), in-depth

Figure 2: Architecture of the Vision Transformer as proposed in [15] and the detailed structure of the Vision Transformer encoder block. In the Vision Transformer, sequential image patches are used as the input and processed using a Transformer encoder to produce the final classification output.

analysis (e.g. task-specific research progress and limitations), as well as a discussion of various aspects.

- Finally, We discuss challenges and open issues and also identify new trends, raise open questions and identify future directions.

***Search Strategy***. We conducted a thorough search using DBLP, Google Scholar, and Arxiv Sanity Preserver, utilizing customized search queries that allowed us to obtain lists of scholarly publications. These publications included peer-reviewed journal papers, conference or workshop papers, non-peer-reviewed papers, and preprints. Our search queries consisted of keywords (`transformer* | deep* | medical* | {Task}*`), (`transformer | medical*`), (`transformer* | medical* | image* | model*`), (`convolution* | vision* | transformer* | medical*`), where **{Task}** refers to one application covered in this review (see Figure 1). To ensure the selection of relevant papers, we meticulously evaluated their novelty, contribution, and significance, and prioritized those that were first of their kind in the field of medical imaging. Following these criteria, we chose papers with the highest rankings for further examination. It is worth noting that our review may have excluded other significant papers in the field, but our goal was to provide a comprehensive overview of the most important and impactful ones.

***Paper Organizations***. The remaining sections of the paper are organized as follows. In Section 2, we provide an overview of the key components of the well-established Transformer architecture and its clinical importance. Moreover, this section clarifies the categorization of neural network variants in terms of the position where the Transformer is located. Section 3 to Section 9 comprehensively review the applications of Transformers in diverse medical imaging tasks as depicted in Figure 1. For each task, we propose a taxonomy to characterize technical innovations and major use cases. Section 10 presents open challenges and future perspectives of the field as a whole, while finally, Section 11 concludes this work.

## 2. Background

In this section, we first provide an overview of the Transformer module and the key ideas behind its feasible design. Then, we outline a general taxonomy of Transformer-based models, characterized by their core techniques of using Transformers, i.e., whether they are purely Transformer-based, or whether the Transformer module is either used in the encoder, decoder, bottleneck, or skip connection, respectively.

### 2.1. Transformers

The original Transformer [14] was first applied to the task for machine translation as a new attention-driven building block. The vanilla Transformer consists of an encoder and a decoder, each of which is a stack of $L$ tandem of consecutive identical blocks. The Transformer module is convolutional-free and

3

solely based on the self-attention mechanism or attention mechanism in short. Specifically, these attention blocks are neural network layers that relate different positions of a single sequence to compute the sequence's representation. Since the establishment of Transformer models, they have attained remarkable performance in diverse natural language processing tasks [24]. Inspired by this, Dosovitskiy et al. proposed the Vision Transformer (ViT) [15] model as illustrated in Figure 2. When trained on large datasets, for instance, JFT-300M, ViT outperforms the then state-of-the-art, namely ResNet-based models like BiT [25]. In their approach, an image is turned into fixed-sized patches before being flattened into vectors. These vectors are then passed through a trainable linear projection layer that maps them into $N$ vectors with the dimensionality of $D \times N$ is the number of patches. The outputs of this stage are referred to as patch embeddings. To preserve the positional information present within each patch, they add positional embeddings to the patch embeddings. In addition to this, a trainable class embedding is also appended to the patch embeddings before going through the Transformer encoder. The Transformer encoder is comprised of multiple Transformer encoder blocks. There is one multi-head self-attention (MSA) block and an MLP block in each Transformer encoder block. The activations are first normalized using LayerNorm (LN) before going into these blocks in the Transformer encoder block. Furthermore, there are skip connections before the LN that add a copy of these activations to the corresponding MSA or MLP block outputs. In the end, there is an MLP block used as a classification head that maps the output to class predictions. The self-attention mechanism is a key defining characteristic of Transformer models. Hence, we start by introducing the core principle of the attention mechanism.

### 2.1.1. Self-Attention

In a self-attention layer (Figure 3a), the input vector is firstly transformed into three separate vectors, i.e., the query vector $q$, the key vector $k$, and the value vector $v$ with a fixed dimension. These vectors are then packed together into three different weight matrices, namely $W^Q$, $W^K$, and $W^V$. A common form of $Q$, $K$, and $V$ can be formulated as Equation (1) for an input $X$

$$K = W^K X, Q = W^Q X, V = W^V X, \quad (1)$$

where $W^K$, $W^Q$, and $W^V$ refers to the learnable parameters. The scaled dot-product attention mechanism is then formulated as

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (2)$$

where $\sqrt{d_k}$ is a scaling factor, and a softmax operation is applied to the generated attention weights to translate them into a normalized distribution.

### 2.1.2. Multi-Head Self-Attention

The multi-head self-attention (MHSA) mechanism (Figure 3b) has been proposed [14] to model the complex relationships of token entities from different aspects. Specifically, the MHSA block helps the model to jointly attend to information

from multiple representation sub-spaces, as the modeling capability of the single-head attention block is quite coarse. The process of MHSA can be formulated as

$$\text{MultiHead}(Q, K, V) = [Concat(\text{head}_1, \ldots, \text{head}_h)]W^O, \quad (3)$$

where $\text{head}_i = \text{Attention}\left(QW_i^Q, KW_i^K, VW_i^V\right)$, and $W^O$ indicates a linear mapping function to combine multi-head representation. Note that $h$ is a hyper-parameter set to $h = 8$ in the original paper.



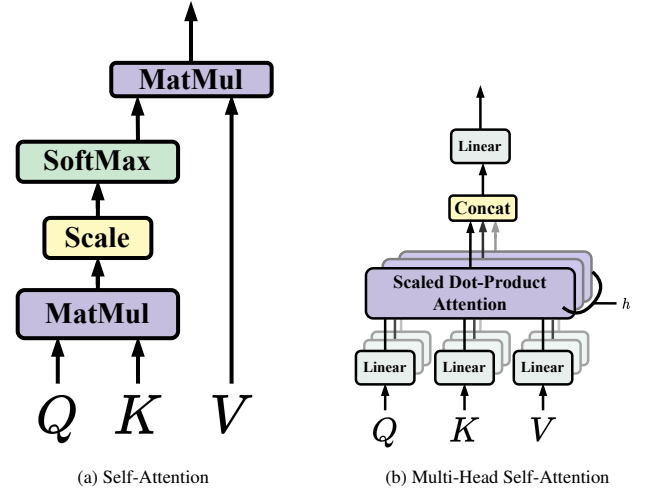(a) Self-Attention    (b) Multi-Head Self-Attention

Figure 3: (a) The process of self-attention. (b) Multi-head attention. The MSA consists of multiple SA blocks (heads) concatenated together channel-wise as proposed in [14].

### 2.2. Transformer modes

While the Transformer was originally introduced with an encoder-decoder pipeline, many modern architectures generally exploit the Transformer architecture in different fashions, which generally depend on the target application. The usage of Transformers in vision tasks can broadly be classified into pure and hybrid designs.

### 2.2.1. Pure Transformers

Due to the deficiency of CNN-based architectures in learning global and long-range semantic information interactions, which stems from the locality of convolution operation, a cohort study has investigated the purely Transformer-based models without any convolution layer. These models usually consist of encoder, bottleneck, decoder, and skip connections directly built upon the ViT or its variants. In this criteria, there are usually multiple multi-head self-attention modules in both encoding and decoding sections that allow the decoder to utilize information from the encoder. Examples of such methods are the Swin-Unet [26] and the TransDeepLab [27] networks which, as their name suggests, try to model the seminal U-Net [28], and DeepLab [29] architectures.

### 2.2.2. Transformer: Hybrid

The hybrid Transformer models usually modify the base CNN structure by replacing the encoder or decoder modules.

**Encoder**: Encoder-only models such as the seminal BERT [30] are designed to make a single prediction per input or a single prediction for an entire input sequence. In the computer vision era, these models are applicable for classification tasks. Moreover, as utilizing a pure Transformer can result in limited localization capacity stemming from inadequate low-level features, many cohort studies try to combine CNN and Transformer in the encoding section [17]. Such a design can enhance finer details by recovering localized spatial information.

**Decoder:** Transformers can also be used in a decoding fashion. Such a causal model is typically used for generation tasks such as language modeling. Besides that, the modification can apply to the skip connections of the decoder module. Skip connection is a widely used technique to improve the performance and the convergence of deep neural networks. It can also serve as a modulating mechanism between the encoder and the decoder. To effectively provide low-level spatial information for the decoding path, the idea of exploiting Transformers in designing skip connections has emerged. This notable idea can lead to finer feature fusion and recalibration while guaranteeing the aggregation scheme of using both high-level and low-level features [17, 31].

### 2.3. Clinical Importance

In practice, medical professionals conduct medical image analysis qualitatively in real-world scenarios. This can lead to different understandings and levels of precision due to variations in the reader's expertise or differences in image quality, as well as being timely and labor-expensive. Therefore, deep learning techniques have gained extensive attraction in medical image analysis aiming to reduce inter-reader variation and decrease the expenses associated with time and workforce [22].

The fundamental question that emerges is **What are the motivations behind using Transformers in the medical domain?**

Advances in adversarial attacks create an unavoidable danger in which potential attackers might attempt to gain benefits by exploiting vulnerabilities in the healthcare system. While there is substantial literature available concerning the robustness of CNNs in the field of medical imaging, it is still a challenging direction to explore. Recent studies [32] prove that ViTs are more robust to adversarial attacks than CNNs. Specifically, ViT is significantly more robust than CNN in a wide range of white-box attacks. A similar trend is also observed in the query-based and transfer-based black-box attacks which can be attributed to the fact that ViTs are more reliant on low-frequency (robust) features while CNNs are more sensitive to high-frequency features.

Besides, due to strict privacy rules, low occurrence rates of certain diseases, concerns about data ownership, and a limited number of patients, providing suitable data have always been a paramount concern in the medical domain. Recently, the idea of exploiting the inherent structure of ViT in distributed medical imaging applications has emerged. Owning to the inherent structure of ViT (as opposed to CNN), it can be split into shared body and task-specific heads which demonstrate that a ViT body with sufficient capacity can be shared across relevant

tasks by leveraging different strategies such as multi-task learning [33]. This way, ViTs can be a paramount application in decentralized medical imaging solutions.

Moreover, Transformers excel at capturing long-range dependencies in data, making them well-suited for analyzing complex medical images where contextual information plays a crucial role. Transformers offer greater interpretability than CNNs, as they provide attention maps that highlight the regions of an image that contribute most to the model's decision-making. This transparency is crucial in medical settings where understanding the reasoning behind a diagnosis is important. The flexible design of ViTs can be used on edge devices to speed diagnosis, improve clinician workflows, enhance patient privacy, and save time. Below, we provide recent examples of clinical use cases of Transformers in the medical domain.

**Med-PaLM 2** [34], a Transformer-based model, by *Google* demonstrated a feasible comparison between physicians and the application of Transformers in the medical domain questioning. The answers provided by Med-PaLM 2 are pretty impressive, and the authors provided a metric to assess model performance by exposing the clinician answers and the model's output to other physicians and lay-persons to rate them in nine categories such as factuality, medical reasoning capability, etc. As a result, the quantitative results endorse that 72.9% of the time, answers provided by Med-PaLM 2 are preferable regarding reporting the medical consensus.

Seenivasan et al. [35] addressed the knowledge transfer's steep learning curve from surgical experts to medical trainees and patients and proposed the **SurgicalGPT**, a Transformer-based multi-modal visual question-answering framework that utilizes the large language models with visual cues. Surgical-GPT demonstrated outperforming results over several (robotic) surgical datasets in terms of accuracy compared to the uni-modality text generation models.

Consequently, the clinical use cases of Med-PaLM 2 and SurgicalGPT have demonstrated the feasibility and benefits of employing Transformers in the medical domain, showing their effectiveness and potential applications.

## 3. Medical Image Classification

Image classification is still one of the challenging problems in computer vision, which aids in segregating extensive quantities of data into meaningful categories [57]. Vision Transformers (ViT) have recently demonstrated outstanding results in various image classification tasks and offer significant advantages over conventional CNNs [58, 59, 60, 61, 62, 63]. These advantages include long-range relationships, adaptive modeling, and attention maps that yield intuition on what the model deems more important inside an image [39]. Due to these alluring advantages, there is rising interest in building Transformer-based models for medical image classification. Therefore, highly precise classification is becoming increasingly vital for facilitating clinical care.

In this section, we exhaustively examine ViTs in medical image classification. As illustrated in Figure 4, we have broadly classified these methods based on the role ViT plays in their
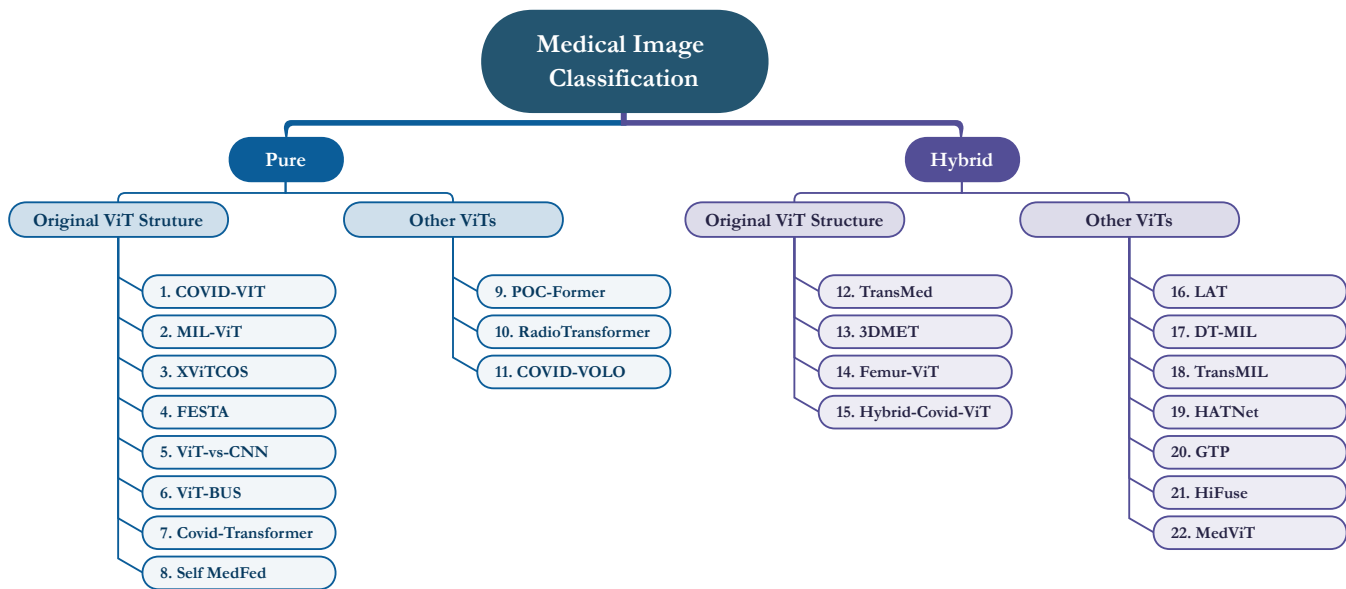
Figure 4: Taxonomy of ViT-based approaches in medical image classification. Methods are categorized based on their proposed architecture into pure and hybrid methods, in which they adopt the vanilla ViT or present a new type of vision Transformer in medical image classification. Notably, we utilize the prefix numbers in the paper's name in ascending order and denote the reference for each study as follows: 1. [36], 2. [37], 3. [38], 4. [33], 5. [39], 6. [40], 7. [41], 8. [42], 9. [43], 10. [44], 11. [45], 12. [46], 13. [47], 14. [48], 15. [49], 16. [50], 17. [51], 18. [52], 19. [53], 20. [54], 21. [55], 22.[56].

architecture. These categories include pure Transformers and Hybrid Models. Generally, a vision Transformer-based classification architecture consists of three modules: (1) a backbone for capturing input features, (2) an encoder for modeling the information, and (3) a classification head for generating output based on the specified task. Therefore, the Transformer can be adopted in each module. However, some works, including Lesion Aware Transformer (LAT) [50] and Deformable Transformer for Multi-Instance Learning (DT-MIL) [51], take a different approach and utilize encoder-decoder structures. LAT proposes a unified encoder-decoder system for Diabetic Retinopathy (DR) grading, and DT-MIL introduces a Transformer-based encoder-decoder architecture for classifying histopathological images, where the deformable Transformer was embraced for the encoder part. In the following, we will go into great depth on both hybrid and pure models.

### 3.1. Pure Transformers

Since the emergence of Transformers, there has been a growing debate regarding whether it is time to switch entirely from CNNs to Transformers. Matsoukas et al. [39] conduct a series of experiments to answer this critical question. They take ResNet50 [64] and the DeiT-S [65] models to represent CNN and ViT models, respectively. They train each of these two models in 3 different fashions: a) randomly initialized weights, b) pre-trained on ImageNet (transfer learning), and c) pre-training on the target dataset in a self-supervised learning (SSL) scheme using DINO [66]. Their findings show that when utilizing random initialization, ViTs are inferior to CNNs. In the case of transfer learning, the results are similar for both models, with ViT being superior for two out of three datasets. Additionally, ViT performs better when self-supervision on the target data is applied. They conclude that Vision Transformers, indeed, are suitable replacements for CNNs.

Transformers have had a profound effect on medical development. Researchers have thoroughly investigated adopting the ViT in medical image classification tasks since its introduction. However, the limited number of medical images has hindered Transformers from replicating their success in medical image classification. **ViT-BUS** [40] studies the use of ViTs in medical ultrasound (US) image classification for the first time. They propose to transfer pre-trained ViT models based on the breast US dataset to compensate for the data-hunger of ViTs. Evaluated results on B [67], BUSI [68], and B+BUSI datasets indicate the predominance of attention-based ViT models over CNNs on US datasets. Likewise, **COVID-Transformer** [41] utilizes ViT-L/16 to detect COVID from Non-COVID based on CXR images. Due to the limitation of sufficient data, they introduce a balanced dataset containing 30K chest X-ray images for multi-class classification and 20K images for binary classification. The published dataset is created by merging datasets [69], [70], and [71]. They fine-tune the model on the dataset with a custom MLP block on top of ViT to classify chest X-ray (CXR) images. Moreover, COVID-Transformer exploits the GradCAM Map [72] to visualize affected lung areas that are significant for disease prediction and progression to display the model interpretability. Similarly, Mondal et al. [38] present **xViTCOS** for detecting COVID-19 in CTs and CXRs. xViT-COS employs a model that has been pre-trained on ImageNet-21k [73]. Nevertheless, the training data capacity might overshadow the generalization performance of the pre-trained ViT to transfer the knowledge from the learned domain to the target domain. By training the model on the COVIDx-CT-2A dataset [74], a moderately-sized dataset, xViTCOS overcomes this problem. However, due to the shortage of the insufficient amount of CXR images, the pre-trained ViT model is fine-tuned using the CheXpert dataset [75]. In addition, xViTCOS

leverages the Gradient Attention Rollout algorithm [76] to visually demonstrate the model's prediction on the input image for clinically interpretable and explainable visualization. In experiments using COVID CT-2A and their custom-collected Chest X-ray dataset, xViTCOS significantly outperforms conventional COVID-19 detection approaches. **MIL-VT** [37] similarly suggests pre-training the Transformer on a fundus image large dataset beforehand, initialized by the pre-trained weight of ImageNet, then fine-tuning it on the downstream retinal disease classification task in order to encourage the model to learn global information and achieve generalization. Unlike previous approaches, they apply some modifications to the vanilla ViT structure. In the classic ViT, embedded features are neglected for classification; instead, only the class token, which retains the summarization of embedded features' information, is used. Yu et al. [37] propose a novel multiple-instance learning (MIL)-head module to exploit those embedded features to complement the class token. This head comprises three submodules that attach to the ViT in a plug-and-play manner: 1) the MIL embedding submodule that maps the feature embeddings to a low-dimensional embedding vector, 2) the attention aggregation submodule that outputs a spatial weight matrix for the low-dimensional patch embeddings; this weight matrix is then applied to the low-dimensional embeddings to ascertain each instance's importance, 3) the MIL classifier submodule that determines the probability of each class through aggregated features. In the downstream task, both MLP and MIL heads use the weighted cross-entropy loss function for training. The outputs of both heads are then weight-averaged for the inference time. Results indicate the effectiveness of the proposed training strategy and the MIL-head module by dramatically boosting the performance over APTOS2019 [77] and RFMiD2020 [78] datasets when compared to CNN-based baselines. In contrast to the previous 2D-based methods that employ transfer learning, **COVID-ViT** [36] proposes training ViT to classify COVID and non-COVID cases using 3D CT lung images. Given that a COVID volume may contain non-COVID 2D slices, COVID-ViT applies a slice voting mechanism after the ViT classification result in which the subject is categorized as having COVID if more than a certain percentage of slices (e.g., 25%) are predicted to be COVID. The findings reported for the MIA-COVID19 competition [79] confirm that ViT outperforms CNN-based approaches such as DenseNet [80] in identifying COVID from CT images.

Besides the remarkable accuracy of Transformers compared to CNNs, one of their major drawbacks is their high computational cost, thereby making them less effective for real-world applications, such as detecting COVID-19 in real time. In light of the prevalence of COVID-19, rapid diagnosis will be beneficial for starting the proper course of medical treatment. CXR and lung CT scans are the most common imaging techniques employed. However, CT imaging is a time-consuming process, and using CXR images is unreliable in identifying COVID-19 in the early stage. In addition, vision Transformers are computationally expensive to deploy on mobile devices for real-time COVID-19 classification. Therefore, Perera et al. [43] present a lightweight **Point-of-Care Transformer (POCFormer)**. The compactness of POCFormer allows for real-time diagnosis of COVID-19 utilizing commercially accessible POC ultrasound devices. POCFormer reduces the complexity of the vanilla ViT self-attention mechanism from quadratic to linear using Linformer [81]. The results display the superiority of POCFormer in the real-time detection of COVID-19 over the CNN-based SOTAs on the POCUS dataset [82].

In addition, despite the great potential shown by ViTs in ImageNet classification, their performance is still lower than the latest SOTA CNNs without additional data. These Transformers mainly focus on a coarse level by adopting a self-attention mechanism to establish global dependency between input tokens. However, relying only on a coarse level restricts the Transformer's ability to achieve higher performance. Thus, Liu et al. [45] leverage a pre-trained version of **VOLO** for an X-ray COVID-19 classification. VOLO [83] first encodes fine-level information into the token representations through proposed outlook attention, alleviating the limitations of Transformers that require a large amount of data for training, and second aggregates the global features via self-attention at the coarse level. Through the outlook attention mechanism, VOLO dynamically combines fine-level features by treating each spatial coordinate $(i, j)$ as the center of a $K \times K$ local window and calculating its similarity with all its neighbors. The findings indicate that fine-tuning VOLO on Dataset-1 [84] leads to 99.67% top1 accuracy on Dataset-1 test cases and 98.98% top1 accuracy on unseen Dataset-2 [85], which demonstrates the generality of the approach.

Furthermore, accessible labeled images have considerably influenced research on the use of Transformers to diagnose COVID-19. Considering the shortage of labeled data, data sharing between hospitals is needed so as to create a viable centralized dataset. However, such collaboration is challenging due to privacy concerns and patient permission. Motivated by Federated Learning (FL) and Split Learning (SL), Park et al. [33] present a **Federated Split Task-Agnostic (FESTA)** framework that uses ViT for multi-task learning of classification, detection, and segmentation of COVID-19 CXR images. FESTA benefits from the decomposable modular design of ViT to train heads and tails via clients and share the server-side Transformer body across clients to aggregate extracted features and process each task. The embedded features from the body Transformer are then passed to their task-specific tail on the client side to produce the final prediction (Figure 5(a)). Figure 5(b) illustrates the single-task learning scheme and (c) the multi-task learning scheme. In multi-task learning, heads, tails, and a task-agnostic Transformer body are first jointly trained for 6000 rounds (see Figure 5(c)). Then, heads and tails are fine-tuned according to the desired specific task while freezing the weights of the Transformer body. FESTA merits from 220000 decentralized CXR images and attains competitive results compared to the data-centralized training approaches. The experimental results also demonstrate the stable generalization performance of FESTA, where multi-task learning enhances the performance of the individual tasks through their mutual effect during training.

Most attention-based networks utilized for detection and classification rely on the neural network to learn the neces-
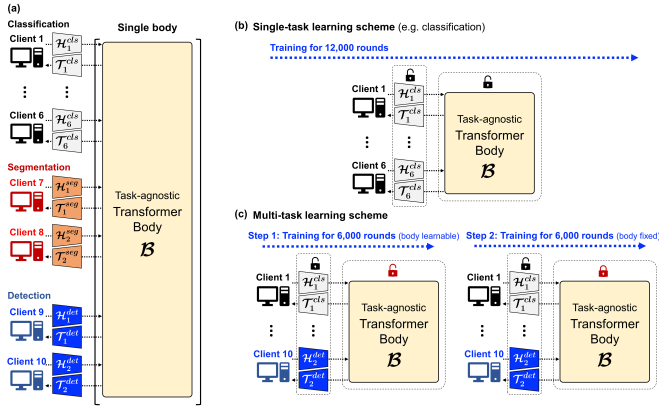
Figure 5: Overview of the FESTA framework [33], which utilizes ViT for multi-task learning of COVID-19 CXR classification, detection, and segmentation. (a) FESTA leverages ViT's decomposable modular design to train heads ($\mathcal{H}$) and tails ($\mathcal{T}$) via clients while sharing the server-side Transformer body ($\mathcal{B}$) between clients to integrate retrieved features. Final predictions are then derived by feeding embedded features to their task-specific tails on the client side. (b) illustrates the single-task learning scheme, and (c) two steps multi-task learning scheme. The former is trained for 12000 rounds, while the latter undergoes two training steps. First, the whole parts train in 6000 rounds. Then by freezing the weights of the Transformer body, the heads and tails are fine-tuned for 6000 steps based on the desired specific task.



Figure 6: Overview of RadioTransformer [44]. Human Visual Attention Training (HVAT) block first uses radiologists' visual observations of chest radiographs to train a global-focal teacher network. The pre-trained teacher network is then utilized to distill the teacher's knowledge to a global-focal student network through visual attention loss, enabling the student to learn visual information. Following the teacher-student strategy and incorporating radiologist visual examinations leads to an improvement in the classification of disease on chest radiographs.

sary regions of interest. Bhattacharya et al. [44] in **Radio-Transformer** argue that in certain applications, utilizing experts' opinions can prove beneficial. Specifically, they apply this notion to leverage radiologists' gaze patterns while diagnosing different diseases on medical images; then, using a teacher-student architecture, they teach a model to pay attention to regions of an image that a specialist is most likely to examine. As illustrated in Figure 6, The teacher and the student networks consist of two main components: global and focal. The global component learns coarse representation while the focal module works on low-level features, and both these segments are comprised of Transformer blocks with shifting windows. In addition, the global and focal components are interconnected using two-way lateral connections to form the global-focal module; this is to address the inherent attention gap between the two. The teacher network is first directly pre-trained on human visual attention maps. Then, the entire model is trained for different downstream tasks, e.g., object detection and classification. Furthermore, the authors propose a self-supervised Visual Attention Loss (VAL) that incorporates both GIoU and MSE loss. The student network is trained to predict probability values for different classes and attention regions. These attention regions are then compared to those obtained from the teacher model, and the weights are optimized using VAL.

### 3.2. Hybrid Models

In spite of the vision Transformers' ability to model global contextual representations, the self-attention mechanism undermines the representation of low-level details. CNN-Transformer hybrid approaches have been proposed to ameliorate the problem above by encoding both global and local features using the locality of CNNs and the long-range dependency of Transformers.
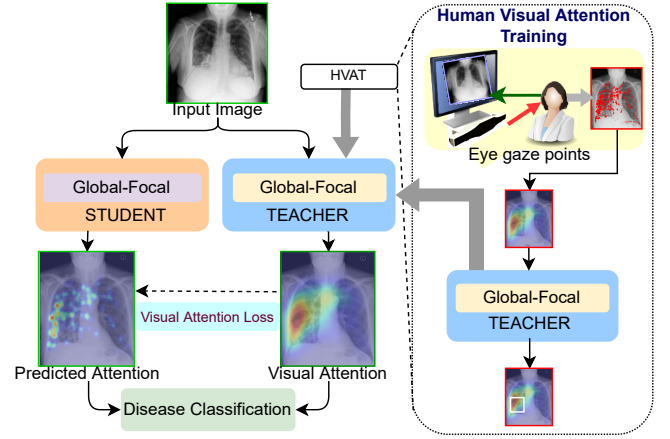
**TransMed** [46] proposes a hybrid CNN-Transformer network that leverages the locality of CNNs and the long-range dependency character of Transformers for parotid gland tumor and knee injury classification. Multimodal medical images primarily have long-range interdependencies, and improving performance requires an effective fusion strategy. TransMed proposes a novel image fusion strategy. Firstly, three neighboring 2D slices of a multimodal image are overlaid to create three-channel images. Then, each image is partitioned into $K \times K$ patches. This fusion approach allows the following network to learn mutual information from images of different modalities. Patch tokens are fed into a CNN network to capture their low-level features and generate patch embeddings. The classic ViT is then used to determine the relationship between patch sequences. TransMed's final results verify the effectiveness of hybrid models in classifying multimodal medical images by outperforming all its counterparts by a large margin. TransMed-S enhances average accuracy on the PGT dataset by about 10.1% over its nearest counterpart, BoTNet [87], while requiring fewer parameters and FLOP count. Comparably, Tanzi et al. [48] develop a new CAD system (**Femur-ViT**) based on Vision Transformers for diagnosing femoral fractures. First, YOLOv3 [88] is utilized to detect and crop the left and right femur regions. Afterward, a CNN (InceptionV3 [89]) and a hierarchical CNN (different InceptionV3 networks in cascade) [90] are applied to the dataset, and the results serve as baselines for the classification. Then, they use a modified ViT to classify seven different fracture types. Finally, a clustering approach is proposed as an evaluation technique for the ViT encoder. This study highlights the power of using ViT models for medical image classification and the ability of the proposed CAD system to significantly increase clinicians' diagnostic accuracy. **3DMeT** [47] proposes applying a 3D medical image Transformer for assessing knee cartilage defects in three grades: grade 0 (no defect), grade 1 (mild defect), and grade 2 (severe defect). Primar-
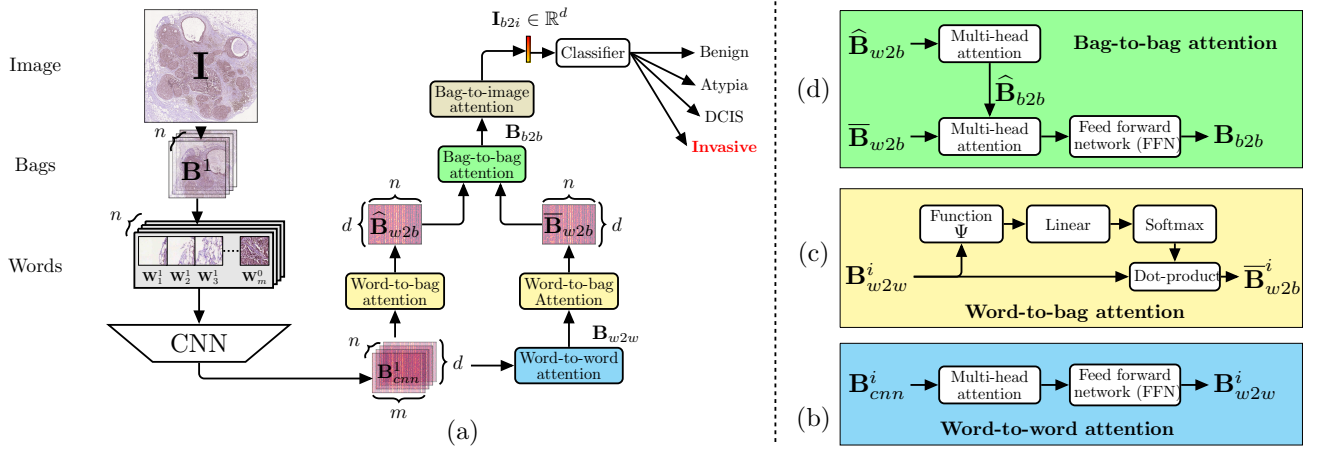
Figure 7: The overall architecture of [86]. HATNet hierarchically divides an input image into $n \times m$ words, which are then fed into the CNN encoder to provide word-level representations for each bag. Then by performing a bottom-up decoding strategy and applying a linear classifier, breast biopsy classification results are obtained. Notably, bag-to-image attention has the same procedure as word-to-bag attention, shown in (c).

ily, using medical 3D volumes as an input to the Transformer is computationally expensive, thereby making it impractical. 3DMeT resolves the high computational cost problem by replacing conventional linear embedding with 3D convolutional layers. The weights of convolutional layers are adopted using the teacher-student training strategy. 3DMeT takes an exponential moving average from the first one/few-layer(s) of the CNN teacher's weights and uses it as convolutional layers' weights. This method enables Transformers to be compatible with small medical datasets and to benefit from CNNs' spatial inductive biases. Lastly, the Transformer and CNN teacher's outputs are combined in order to derive the classification results.

Operating Transformers over Whole Slide Images (WSIs) is computationally challenging since WSI is a gigapixel image that retains the original structure of the tissue. MIL and CNN backbones have demonstrated practical tools for acting on WSI. MIL is a weakly supervised learning approach that enables deep learning methods to train high-resolution images like WSI. Since annotating such images at the pixel level is impractical, MIL proposes to divide an input WSI into a bag of instances and assign a single label to the bag of each image based on pathology diagnosis. The bag has a positive label if it contains at least one positive instance, and it is considered negative if all the instances in the bag are negative. Then CNN backbones are employed to down-sample and extract the features of each instance and allow Transformers to operate according to the generated feature maps and currently available hardware. Therefore, **DT-MIL** [51] proposes to compress WSIs into compact feature images by embedding each patch of the original WSI into a super-pixel at its corresponding position using EfficientNet-B0 [91]. The resulting thumbnail image feed into a $1 \times 1$ Conv for feature reduction, followed by a deformable Transformer encoder that aggregates instance representations globally. A similar approach is adopted by **H**olistic **AT**tention **Net**work (**HATNet**) [53], where they first divide an input image into $n$ non-overlapping bags, each broken down into $m$ non-overlapping words (or patches). $n \times m$ words are fed into the CNN encoder to obtain word-level representations

for each bag. HATNet aims to develop a computer-aided diagnosis system to help pathologists in reducing breast cancer detection errors. According to the World Health Organization (WHO), breast cancer is the most frequent non-skin cancer in women, accounting for one out of every four new female cancers annually [92]. As illustrated in Figure 7, HATNet follows a bottom-up decoding strategy such that it first performs multi-head attention to words in a *word-to-word attention* block, then considers the relationship between words and bags in *word-to-bag attention*, followed by *bag-to-bag attention* to attain interbag representations. The acquired bag features are then aggregated in *bag-to-image attention* to build image-level representations. A linear classifier is ultimately applied to achieve the final results. Furthermore, unlike most MIL methods that take all the instances in each bag independent and identically distributed [93, 94, 95], **TransMIL** [52] suggests that it is essential to consider the correlation between different instances and explore both morphological and spatial information. Two Transformer layers address the morphological information, and a conditional position encoding layer named Pyramid Position Encoding Generator (PPEG) addresses the spatial information. The proposed PPEG module has two merits: 1) It handles positional encoding of sequences with a variant number of instances by using group convolution over the 2D reshaped patch tokens, and 2) It enriches the features of tokens by capturing more context information through convolutions. In contrast to conventional IID-based MIL methods requiring many epochs to converge, TransMIL converges two to three times faster by using morphological and spatial information. TransMIL also outperforms all the latest MIL methods [96, 97, 98, 99, 93] in terms of accuracy and AUC by a significant margin in binary and multiple classification tasks and exhibits the superiority of taking the correlation between different instances into account and considering both morphological and spatial information.

Previous methods mainly rely on weakly supervised learning or dividing WSIs into image patches and using supervised learning to assess the overall disease grade. Nevertheless, these approaches overlook WSI contextual information. Thus, Zheng

9

Table 1: An overview of the reviewed Transformer-based medical image classification models.

| Method | Modality | Organ | Type | Pre-trained Module: Type | Datasets | Metrics | Year |
|--------|----------|-------|------|--------------------------|----------|---------|------|
| **Pure** | | | | | | | |
| ViT-vs-CNN [39] | Fundus Dermoscopy X-ray | Eye Skin Breast | 2D | ViT: Self-supervised & Supervised | [1]APTOS-2019 [77] [2]ISIC-2019 [100] [3]CBIS-DDSM [101] | Kappa Recall ROC-AUC | 2021 |
| ViT-BUS [40] | Ultrasound | Breast | 2D | ViT: Supervised | [1]B [67] [2]BUSI [68] | ACC AUC | 2022 |
| POCFormer [43] | Ultrasound | Chest | 2D | ✗ | POCUS [82] | Recall, F1 SP, SE, ACC | 2021 |
| MIL-VT [37] | Fundus | Eye | 2D | ViT: Supervised | [1]Private Dataset [2]APTOS-2019 [77] [3]RFMiD-2020 [78] | Recall, F1 ACC, AUC Precision | 2021 |
| COVID-VIT [36] | CT | Chest | 3D | ✗ | MIA-COV19 [79] | ACC, F1 | 2021 |
| xViTCOS [38] | X-ray CT | Chest | 2D | ViT: Supervised | [1]COVID CT-2A [74] [2]CheXpert [75] | Recall, F1 Precision SP, NPV | 2021 |
| FESTA [33] | X-ray | Chest | 2D | ViT: Supervised | [1]Four Private Datasets [2]CheXpert [75], [3]BIMCV [102] [4]Brixia [103], [5]NIH [104] [6]SIIM-ACR [105], [7]RSNA [106] | Recall, F1 SP, SE, AUC | 2021 |
| COVID-Transformer [41] | X-ray | Chest | 2D | ViT: Supervised | [1][69], [2][71] [3][70] | Recall, F1 ACC, AUC Precision | 2021 |
| COVID-VOLO [45] | X-ray | Chest | 2D | ViT: Supervised | [1][84] [2][85] | ACC | 2021 |
| RadioTransformer [44] | X-ray | Chest | 2D | ViT: Supervised | [1]RSNA [107], [2]Cell Pneumonia [108] [3]COVID-19 Radiography [84, 109] [4]NIH [104], [5]VinBigData [110] [6]SIIM-FISABIO-RSNA [111] [7]RSNA-MIDRC [112, 113] [8]TCIA-SBU COVID-19 [114, 115] | Recall, F1 ACC, AUC Precision | 2022 |
| Self MedFed [42] | X-ray Dermoscopy Fundus | Chest Skin Eye | 2D | ViT: Self-supervised | [1]EyePACKS [116] [2]ISIC 2017 [117], [3]ISIC 2018 [118] [4]ISIC 2020 [119] [5]COVID-FL [42] | ACC | 2023 |
| **Hybrid** | | | | | | | |
| TransMIL [52] | Microscopy | Multi-organ | 2D | CNN: Supervised | [1]Camelyon16 [120] [2]TCGA-NSCLC [121, 122] [3]TCGA-RCC [123] | ACC AUC | 2021 |
| LAT [50] | Fundus | Eye | 2D | CNN: Supervised | [1]Messidor-1 [124] [2]Messidor-2 [125] [3]EyePACKS [116] | AUC & Kappa | 2021 |
| TransMed [46] | MRI | Ear Knee | 3D | ViT: Supervised CNN: Supervised | [1]PGT [46] [2]MRNET [126] | Precision ACC | 2021 |
| 3DMeT [47] | MRI | Knee | 3D | CNN: Supervised | Private dataset | ACC, Recall, F1 | 2021 |
| Hybrid-COVID-ViT [49] | X-ray | Chest | 2D | CNN: Supervised | CheXpert [75] | AUC, ACC SP, SE | 2021 |
| Femur-ViT [48] | X-ray | Femur | 2D | ViT: Supervised CNN: Unsupervised | Private dataset | Recall, F1 Precision, ACC | 2022 |
| DT-MIL [51] | Microscopy | Lung Breast | 2D | CNN: Supervised | [1]CPTAC-LUAD [114] [2]BREAST-LNM [51] | Recall, F1 AUC, precision | 2021 |
| GTP [54] | Microscopy | Lung | 2D | CNN: Self-supervised | [1]NLST [127], [2]CPTAC [128] [3]TCGA [129] | Precision, Recall SP, SE, ACC, AUC | 2022 |
| HATNet [53] | Microscopy | Breast | 2D | CNN: Supervised | Breast Biopsy WSI Dataset [130] | ACC, ROC-AUC F1, SP, SE | 2022 |
| HiFuse [55] | CT Dermoscopy Endoscopy | Skin Chest Stomach | 2D | ✗ | [1]ISIC 2018 [118] [2]COVID19-CT [131] [3]Kvasir [132] | Recall, ACC F1, Precision | 2022 |
| MedViT [56] | CT X-ray ultrasound OCT | Multi-organ | 2D | ✗ | MedMNIST [133] | AUC, ACC Top-1 | 2023 |

et al. [54] propose a **Graph-based Vision Transformer (GTP)** framework for predicting disease grade using both morphological and spatial information at the WSIs. The graph term allows for the representation of the entire WSI, and the Transformer term allows for computationally efficient WSI-level analysis. The input WSI is first divided into patches, and those that con-

Table 2: A brief description of the reviewed Transformer-based medical image classification models. The unreported number of parameters indicates that the value was not mentioned in the paper, and the code was unavailable.

| Method | # Params | Contributions | Highlights |
|---|---|---|---|
| **Pure** | | | |
| ViT-vs-CNN [39] | 22M | • They investigate three different weight initialization approaches on three medical datasets: (1) random initialization, (2) transfer learning using supervised ImageNet pre-trained weights, and (3) self-supervised pretraining on the target dataset using DINO [66]. Their final verdict is that ViTs can replace CNNs. | • Utilize three different training schemes on three different datasets to conclude whether ViT can replace CNN. <br>• They repeat their processes five times to be certain of the outcome. <br>• Comparing only two models, DeiT-S and Resnet-50, on only three datasets cannot generalize the conclusion of the superiority of each of the Transformer and CNN. |
| ViT-BUS [40] | ViT-Ti/16: 5.53M <br> ViT-S/32: 22.87M <br> ViT-B/32: 87.44M | • Proposes the use of ViT for the classification of breast ultrasound images for the first time. | • Transferring pre-trained ViT models based on small ultrasound datasets yields much higher accuracy than CNN models. |
| POCFormer [43] | Binary CLS: 2.8M <br> Multiclass CLS: 6.9M | • Proposes a lightweight Transformer architecture that uses lung ultrasound images for real-time detection of COVID-19. | • POCFormer can perform in real-time and be deployed to portable devices. <br>• POCFormer can be used for rapid mass testing of COVID-19 due to its compactness. |
| MIL-VT [37] | 22.12M | • Proposes to first pre-train the Vision Transformer on a fundus image large dataset and then fine-tune it on the downstream task of the retinal disease classification. <br>• Introduces the MIL-VT framework with a novel Multiple Instance Learning(MIL)-head to effectively utilize embedding features to improve the ViT performance. | • The MIL head can significantly enhance the performance by easily attaching to the Transformer in a plug-and-play manner. <br>• MIL-VT efficiently exploits the embedding features overlooked in the final prediction of the ViT. |
| COVID-VIT [36] | 52.81M | • Offers to utilize ViT to classify COVID and non-COVID patients using 3D CT lung images. | • COVID-ViT performs better in classifying COVID from Non-COVID compared to DenseNet. <br>• The reported result is not enough to conclude. They only compare ViT with DenseNet. |
| xViTCOS [38] | 85.99 | • Proposes and explores using ViT for detecting COVID-19 from CXR and CT images. <br>• xViTCOS makes use of the Gradient Attention Rollout algorithm [76] for visualization and clinical interpretability of the output. | • Uses a heatmap plot to demonstrate the model's explainability. |
| FESTA [33] | Body: 66.37M <br> (CLS) Head: 13.31M, Tail: 2k <br> (SEG) Head: 15.04M, Tail: 7.39M <br> (DET) Head: 25.09M, Tail: 19.77M | • Proposes a Federated Split Task-Agnostic (FESTA) framework that leverages ViT to merit from federated learning and split learning. <br>• They use multi-task learning to classify, detect, and segment COVID-19 CXR images. | • The proposed FESTA Transformer improved individual task performance when combined with multi-task learning. <br>• Experimental results demonstrate stable generalization and SOTA performance of FESTA in the external test dataset even under non-independent and non-identically distributed (non-IID) settings. <br>• FESTA eliminates the need for data exchange between health centers while maintaining data privacy. <br>• Using FESTA in the industry may not be safe because it may encounter external attacks on the server that may lose the parameters of the entire network. In addition, using this method may jeopardize patient information through privacy attacks. <br>• Authors did not evaluate the robustness of their approach to difficulties through communication, stragglers, and fault tolerance. |
| COVID-Transformer [41] | ViT-L/16: 307M | • COVID-Transformer investigates using ViT for detecting COVID-19 from CXR images. <br>• COVID-Transformer introduces a new balanced chest X-ray dataset containing 30K images for multi-class classification and 20K for binary classification. | • Uses a heatmap plot to demonstrate the model's explainability. |
| COVID-VOLO [45] | 86.3M | • Proposes fine-tuning the pre-trained VOLO [83] model for COVID-19 diagnosis. | • Using VOLO enables capturing both fine-level and coarse-level features resulting in higher performance in COVID-19 binary classification. |
| RadioTransformer [44] | 3.93M | • Presents a novel global-focal RadioTransformer architecture, including Transformer blocks with shifting windows, to improve diagnosis accuracy by leveraging the knowledge of experts. <br>• Introduces an innovative technique for training student networks by utilizing visual attention regions generated by teacher networks. | • outperforms counterpart backbones on multiple datasets. <br>• Model's explainability |
| Self MedFed [42] | 85.2M | • Self MedFed is a novel privacy-preserving federated self-supervised pre-training framework designed to learn visual representations from decentralized data by leveraging masked image modeling. <br>• Introduces a benchmark dataset called COVID-FL, consisting of federated chest X-ray data from eight different medical datasets. | • The proposed network simultaneously tackles the dual challenges of data heterogeneity and label deficiency. <br>• Experimental results on various medical datasets validate the superiority of the proposed method compared to ImageNet-supervised baselines and existing federated learning algorithms in terms of label efficiency and robustness to non-IID data. |
| **Hybrid** | | | |
| TransMIL [52] | 2.67M | • Presents a Transformer-based Multiple Instance Learning (MIL) approach that uses both morphological and spatial information for weakly supervised WSI classification. <br>• Proposes to consider the correlation between different instances of WSI instead of assuming them independently and identically distributed <br>• Proposes a CNN-based PPEG module for conditional position encoding, which is adaptive to the number of tokens in the corresponding sequence | • Converges two to three times faster than SOTA MIL methods. <br>• The proposed method can be employed for unbalanced/balanced and binary/multiple classification with great visualization and interpretability. <br>• TransMIL is adaptive for positional encoding as token numbers in the sequences changes. <br>• It needs further improvement to handle higher magnification than ×20 of WSIs - Higher magnification means longer sequences, which in turn require more memory and computational costs to process. |
| LAT [50] | - | • Proposes a unified Transformer-based encoder-decoder structure capable of DR grading and lesion detection simultaneously. <br>• Proposes a Transformer-based decoder to formulate lesion discovery as a weakly supervised lesion localization problem. <br>• Proposes lesion region importance mechanism to determine the importance of lesion-aware features. <br>• Proposes lesion region diversity mechanism to diversify and compact lesion-aware features. | • Unlike most approaches that confront lesion discovery and diabetic retinopathy grading tasks independently, which may generate suboptimal results, the proposed encoder-decoder structure is jointly optimized for both tasks. <br>• Despite existing methods that only perform well for discovering explicit lesion regions, LAT can also detect less dense lesion areas. <br>• The proposed LAT is capable of identifying Grades 0 and 1, which are hard to distinguish. |
| TransMed [46] | TransMed-Tu: 17M <br> TransMed-S: 43M <br> TransMed-B: 110M <br> TransMed-L: 145M | • Proposes a hybrid CNN-Transformer network for multimodal medical image classification. <br>• Proposes a novel image fusion strategy for 3D MRI data. | • TransMed achieves much higher accuracy in classifying parotid tumors and knee injuries than CNN models. <br>• Requiring fewer computational resources compared to SOTA CNNs. |
| 3DMeT [47] | - | • Replaces conventional image embedding with 3D convolution layers to reduce the computational cost of using 3D volumes as the Transformer's inputs. <br>• Obtains weights for 3D convolution layers by using a teacher-student training strategy. | • The proposed method makes the Transformers capable of using 3D medical images as input. <br>• 3DMeT Uses significantly fewer computational resources. <br>• Adopting CNN as a teacher assists in inheriting CNN's spatial inductive biases. |
| Hybrid-COVID-ViT [49] | - | • Proposes a vision Transformer that embeds features for high-level COVID-19 diagnosis classification using a backbone trained to spot low-level abnormalities in CXR images. | • Different from SOTA models, the proposed model does not use the ImageNet pre-trained weights while archiving significantly better results. <br>• They examine the interpretability of the proposed model. |
| Femur-ViT [48] | - | • Investigates using ViT for classifying femur fractures. <br>• proposes using unsupervised learning to evaluate the ViT results. | • Achieves SOTA results compared to the CNN models. |
| GTP [54] | - | • Proposes a graph-based vision Transformer (GTP) framework for predicting disease grade using both morphological and spatial information at the WSIs. <br>• Proposes a graph-based class activation mapping (GraphCAM) method that captures regional and contextual information and highlights the class-specific regions. <br>• They use a self-supervised contrastive learning approach to extract more robust and richer patch features. <br>• They exploit a mincut pooling layer [134] before the vision Transformer layer to lessen the number of Transformer input tokens and reduce the model complexity. | • In contrast to SOTA approaches, the proposed GTP can operate on the entire WSI by taking advantage of graph representations. In addition, GTP can efficiently classify disease gade by leveraging a vision Transformer. <br>• The proposed GTP is interpretable so that it can identify salient WSI areas associated with the Transformer output class. <br>• GTP obviates the need for adding extra learnable positional embeddings to nodes by using the graph adjacency matrix. It enables diminishing the complexity of the model. <br>• The proposed GTP takes both morphological and spatial information into account. |
| DT-MIL [51] | 10.88M | • Presents a novel embedded-space MIL approach incorporated with an encoder-decoder Transformer for histopathological image analysis. Encoding is done with a deformable Transformer, and decoding with a classic ViT. <br>• An efficient method to render a huge WSI is proposed, which encodes the WSI into a position-encoded feature image. <br>• The proposed method selects the most discriminative instances simultaneously by utilizing associated attention weights and calibrating instance features using the deformable self-attention. | • The proposed method efficiently embeds instances' position relationships and context information into bag embedding. <br> An extensive analysis of four different bag-embedding modules is presented on two datasets. |
| HATNet [53] | (w/ MobileNetv2): 5.59M <br> (w/ ESPNetv2): 5.58M <br> (w/ MNASNet): 5.47M | • Presents a novel end-to-end hybrid method for classifying histopathological images. | • HATNet surpasses the bag-of-words models by following a bottom-up strategy and taking into account inter-word, word-to-bag, inter-bag, and bag-to-image representations, respectively. (word → bag → image) |
| HiFuse [53] | HiFuse-Ti: 82.49M <br> HiFuse-S: 93.82M <br> HiFuse-B: 127.80M | • A parallel framework is developed to effectively capture local spatial context features and global semantic information representation at various scales. <br>• An adaptive hierarchical feature fusion (HFF) block is created to selectively blend semantic information from various scale features of each branch using spatial attention, channel attention, residual inverted MLP, and shortcut connection. | • The HiFuse model yields favorable outcomes on the three different ISIC 2018 [117], Covid19-CT [131], and Kvasir [132] datasets. <br>• HiFuse incorporates a modular design that offers rich scalability and linear computational complexity. |
| MedViT [56] | MedViT-Ti: 10.2M <br> MedViT-S: 23M <br> MedViT-B: 45M | • Proposes a hybrid model that combines the locality of CNNs and the global connectivity of vision Transformers, aiming to improve the reliability of deep medical diagnosis systems. <br>• By utilizing an efficient convolution operation, the proposed attention mechanism enables joint attention to the information in various representation subspaces, thereby reducing computational complexity while maintaining performance. <br>• Introduces modifying the shape information in the high-level feature space by permuting the feature mean and variance within mini-batches to improve the model's resilience against adversarial attacks. | • The proposed attention mechanism is efficient. <br>• MedViT enhances the robustness of the Transformer models against adversarial attacks. |

tain more than 50% of the background are eliminated and not considered for further processing. Selected patches are fed forward through a contrastive learning-based patch embedding module for feature extraction. A graph is then built via a graph construction module utilizing patch embeddings as nodes of the graph. In the graph Transformer section, a graph convolution

layer followed by a mincut pooling layer [134] is applied first to learn and enrich node embeddings and then lessen the number of Transformer input tokens. Since the graph adjacency matrix contains spatial information of nodes, by adding an adjacency matrix to node features, GTP obviates the need for adding extra learnable positional embeddings to nodes. The final Transformer layer predicts the WSI-level class label for three lung tumor classes: Normal, LUAD, and LSCC. GTP also introduces a graph-based class activation mapping (GraphCAM) technique that highlights the class-specific regions. GraphCAM exploits attention maps from multi-head self-attention (MHSA) blocks in the Transformer layer and maps them to the graph space to create a heatmap for the predicted class. The experiments show that GTP performs as a superior interpretable and efficient framework for classifying WSI images while considering morphological and spatial information.

Diabetic Retinopathy (DR) is an eye disorder that can cause impaired vision and sightlessness by damaging blood vessels in the retina. Most deep-learning approaches view lesion discovery and DR grading as independent tasks that may produce suboptimal results. In contrast to conventional methods, LAT [50] proposes a unified encoder-decoder structure that comprises a pixel relation-based encoder to capture the image context information and a lesion filter-based decoder to discover lesion locations, which the whole network jointly optimized and complemented during training. The encoder is particularly in charge of modeling the pixel correlations, and the Transformer-based decoder part is formulated as a weakly supervised localization problem to detect lesion regions and categories with only DR severity level labels. In addition, LAT proposes two novel mechanisms to improve the effectiveness of lesion-aware filters: 1) Lesion region importance mechanism, $g(\cdot|\Phi)$, to determine the contribution of each lesion-aware feature, and 2) Lesion region diversity mechanism to diversify and compact lesion-aware features. The former is a linear layer followed by a sigmoid activation function that generates importance weights for lesion-aware features, and the latter adopts a triplet loss [135] to encourage lesion filters to find diverse lesion regions. In the DR grading branch, LAT presents a DR grading classification module that calculates a global consistency loss based on the lesion-aware features, indicated as $h(\cdot|\sigma)$. Eventually, the final DR grading prediction is achieved by calculating the cross-entropy loss between the predicted labels obtained from the fusion of $g(\cdot|\Phi)$ and $h(\cdot|\sigma)$ and the ground truth. The total loss is the aggregation of cross-entropy loss, global consistency loss, and triplet loss. Visual results of LAT regarding the lesion discovery are depicted in Figure 8.

The current success of deep learning in medical image interpretation relies heavily on supervised learning, which requires large labeled datasets. However, annotating medical imaging data is costly and time-consuming. Additionally, models struggle to generalize across external institutions or different tasks. Self-supervised learning offers a solution by leveraging unlabeled medical data to develop robust models without costly annotations [137]. Chen et al. [138] propose a valuable empirical study of training self-supervised vision Transformers, examining their effectiveness and performance. Accord-
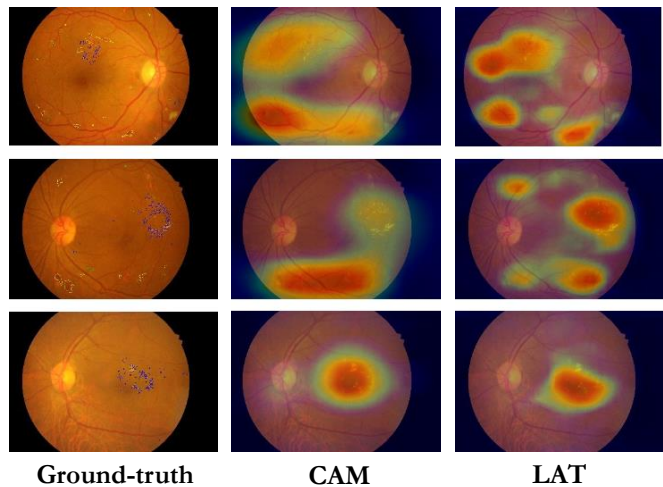


| Ground-truth | CAM | LAT |

Figure 8: LAT [50] vs. CAM [136] visual comparison. The ground truth consists of microaneurysms, hemorrhages, soft exudates, and hard exudates, which are colored as green, yellow, green, and blue dots, respectively (zoom in on the ground truth image for better clarity)

.

ing to their findings, self-supervised Transformers demonstrate strong performance when trained using contrastive learning, outperforming ImageNet-supervised ViT models as their size increases. For example, self-supervised pre-training surpasses supervised pre-training in certain cases, particularly with the very large ViT-Large. Furthermore, self-supervised ViT models achieve competitive results compared to prominent convolutional ResNets in prior studies [139, 140], showcasing the potential of ViT with relatively fewer inductive biases [15]. Notably, they discovered that removing the position embedding in ViT only has a slight negative impact on accuracy, indicating that self-supervised ViT can learn powerful representations without relying heavily on positional information, but also suggesting that the positional information may not have been fully utilized. Therefore, these findings indicate that the combination of self-supervised learning and vision Transformers can be beneficial to adopt for different medical image analysis tasks. For example, Yan et al. [42] propose a robust self-supervised federated learning framework, **Self MedFed**, for medical image classification. Their approach combines masked image modeling with Transformers to pre-train models directly on decentralized target task datasets. This enables more robust representation learning on heterogeneous data and effective knowledge transfer. Experimental results on public medical imaging datasets [116, 117, 118] and their created COVID-FL dataset demonstrate the framework's effectiveness in improving model robustness and learning visual representations across non-IID clients. The research addresses data heterogeneity and label deficiency, showcasing the broad applicability of their self-supervised federated learning framework in medical image analysis.

*3.3. Discussion and Conclusion*

Section 3 thoroughly outlines 22 distinctive Transformer-based models in medical image classification. We have categorized the introduced models based on their architectures

Table 3: Comparison of Transformer-based models on different medical image classification datasets. (†) and (⋆) indicate that results are adopted from [55] and [56].

| ISIC 2018† [118] | | | | | | | TissueMNIST⋆ [141] | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Method | Params (M) | FLOPs (G) | Accuracy % | F1 % | Precision % | Recall % | Method | Params (M) | FLOPs (G) | Top-1% |
| **HiFuse-S [55]** | 93.82 | 8.84 | **83.59** | **72.70** | 72.70 | **73.14** | **MedViT-S [56]** | 23.6 | 4.9 | **73.1** |
| **Conformer-B/16 [142]** | **83.29** | 22.89 | 82.66 | 72.44 | **73.31** | 71.66 | **Twins-SVT-S [143]** | 24.0 | 2.9 | 72.1 |
| **Swin-B [58]** | 87.77 | 15.14 | 79.79 | 63.95 | 65.09 | 63.65 | **PoolFormer-S36 [144]** | 31.2 | 5.0 | 71.8 |
| **ViT-B/32 [15]** | 88.30 | **8.56** | 77.92 | 57.52 | 58.74 | 56.90 | **Swin-T [58]** | 29.0 | 4.5 | 71.7 |
| COVID19-CT† [131] | | | | | | | CvT-13 [145] | 20.1 | 4.5 | 71.6 |
| **HiFuse-S [55]** | 93.82 | 8.84 | **76.88** | **76.31** | **77.78** | 76.19 | **RVT-S [146]** | 22.1 | 4.7 | 71.2 |
| **Conformer-B/16 [142]** | **83.29** | 22.89 | 75.81 | 75.60 | 76.81 | **77.81** | **MedViT-T [56]** | 10.8 | **1.3** | 70.3 |
| **ViT-B/32 [15]** | 88.30 | **8.56** | 61.83 | 60.59 | 61.89 | 60.94 | **MedViT-L [56]** | 45.8 | 13.4 | 69.9 |
| **Swin-B [58]** | 87.77 | 15.14 | 60.75 | 56.36 | 63.20 | 58.95 | **RVT-Ti [146]** | 8.6 | **1.3** | 69.6 |
| Kvasir† [132] | | | | | | | CoaT Tiny [147] | **5.5** | 4.4 | 69.3 |
| **HiFuse-S [55]** | 93.82 | 8.84 | **85.00** | **84.96** | **85.08** | **85.00** | **RVT-B [146]** | 86.2 | 17.7 | 69.3 |
| **Conformer-B/16 [142]** | **83.29** | 22.89 | 84.25 | 84.27 | 84.45 | 84.37 | **DeiT-S [65]** | 22.0 | 4.6 | 67.0 |
| **Swin-B [58]** | 87.77 | 15.14 | 77.30 | 77.29 | 77.74 | 77.44 | **PiT-S [148]** | 23.5 | 2.9 | 66.9 |
| **ViT-B/32 [15]** | 88.30 | **8.56** | 73.80 | 73.50 | 74.24 | 73.72 | **PVT-S [149]** | 25.4 | 4.0 | 66.7 |

into hybrid and pure. These approaches differ according to whether they adhere to the original structure of the vanilla ViT or provide a new variant of the vision Transformer that can be applied to medical applications. In addition, we have presented details on the studied classification methods regarding their architecture type, modality, organ, pre-trained strategy, datasets, metrics, and the year of publication in Table 1. Additional descriptions of the methods, including their model size, contributions, and highlights, are described in Table 2. Table 3 presents a comparison of Transformer-based models across four distinct medical image classification datasets. The evaluation primarily focuses on assessing the accuracies achieved by various models, as well as analyzing their Flop count and number of parameters.

Additionally, time inference and GPU usage are pivotal factors when evaluating method performance. Inference time is a crucial metric for real-time or time-critical applications, and efficient GPU utilization ensures optimal resource allocation, preventing undue stress on the hardware. We conducted experiments using different MedViT models (small, base, and large) on an NVIDIA Tesla T4 GPU, assessing GPU usage and inference times for input images of size $3 \times 224 \times 224$. We performed 1000 model inferences after a warm-up period of 50 iterations to gauge performance. For a batch size of 8, MedViT-S consumes 3.92 GB of GPU memory, while MedViT-B and MedViT-L require 5.04 GB and 6.06 GB, respectively. In terms of inference time (batch size 1), MedViT-S, -B, and -L averaged 49.60 ms, 66.89 ms, and 87.95 ms, respectively. These insights provide valuable information for choosing the right model for specific applications, balancing performance with GPU resource constraints, and enabling informed comparisons with other models when considering factors like the number of parameters and FLOPs. By examining these metrics, researchers can gain insights into the performance and efficiency of different models in the context of medical image classification tasks.

As is evident in the storyline of this section, we have discussed methods in each paragraph regarding the underlying problems in medical image classification and introduced solutions and how they address such issues. However, the need for more research on these problems is crucial to making such approaches widely applicable.

Data availability in the medical domain is one of the most challenging aspects of developing Transformer-based models since Transformer models are known for being data-hungry to generalize. Reasons for data scarcity in the medical field can be referred to as privacy concerns of patients, the time-consuming and costly process of annotation, and the need for expert staff. To this end, the use of generative models [150, 151, 152] and their integration with Transformer models can become prominent since they are capable of creating synthetic data that is comparable to genuine data. In addition, another way to attack this problem is by utilizing federated learning, such as [33]. Nevertheless, there is still room for improvement when it comes to privacy concerns since, in federated learning, communication between the client and server is required.

Despite their SOTA performance, Transformer-based networks still face challenges in deploying their models in the real world due to computational limitations. As shown in Table 2, most approaches have a high number of parameters which provokes a serious problem. Different novel approaches have been introduced to reduce the quadratic complexity of self-attention, which can be leveraged in the medical domain. Furthermore, though ViTs have shown impressive capabilities in ImageNet classification, their performance is still lower than the latest SOTA CNNs without additional data [45]. Hence, existing methods mostly follow pre-training strategies on the ImageNet dataset to build the pre-trained weights for the subsequent downstream tasks. However, despite the enhancement, the domain of natural

images is significantly different from medical data, thereby restricting the performance of further improvement. Therefore, we believe efficient Transformers will considerably influence the future research of Transformer-based models.

## 4. Medical Image Segmentation

Medical segmentation is a significant sub-field of image segmentation in digital image processing. It aims to extract features from a set of regions partitioned from the entire image and segment the key organs simultaneously, which can assist physicians in making an accurate diagnosis in practice. X-ray, positron emission tomography (PET), computed tomography (CT), magnetic resonance imaging (MRI), and ultrasound are common imaging modalities used to collect data. The CNN-based U-Net [28, 4, 166, 167] has been the main choice in this field due to its effective performance and high accuracy. Nevertheless, it cannot extract long-range dependencies in high-dimensional and high-resolution medical images [168, 169]. Therefore, the flexible combination of the U-Net structure with Transformers become a prevalent solution to the segmentation problem at present. Take the multi-organ segmentation task as an example: some networks can achieve state-of-the-art multi-organ segmentation performance on the Synapse dataset (as shown in Figure 10) for abdominal images.

In this section, we present the application of ViTs in segmentation tasks. First, we divide the approaches into two categories: *pure Transformers* and *hybrid Transformers*, where the *pure Transformer* denotes the use of the multiple multi-head self-attention modules in both the encoder and decoder. Hybrid architecture-based approaches fuse the ViTs with convolution modules as the encoder, bottleneck, decoder, or skip connection part to leverage information about the global context and local details. Furthermore, we review some methods with other architectures that propose several novel manners for self-supervised learning. Figure 9 demonstrates the different directions of the methods employing Transformers in the U-Net architecture.

### 4.1. Pure Transformers

In this section, we review several networks referred to as *pure Transformers*, which employ Transformer blocks in both the encoding and the decoding paths. Despite the great success of CNN-based approaches in medical segmentation tasks, these models still have limitations in learning long-range semantic information of medical images. The authors proposed **Swin-Unet**, a symmetric Encoder-Decoder architecture motivated by the hierarchical Swin Transformer [58], to improve segmentation accuracy and robust generalization capability. In contrast to the closest approaches [158, 157, 156, 172] using integrations of CNN with Transformer, Swin-Unet explores the possibility of pure Transformer applied to medical image segmentation.

As shown in Figure 11, Swin-Unet consists of encoder, bottleneck, decoder, and skip connections utilizing the Swin Transformer block with shifted windows as the basic unit. For the encoder, the sequence of embeddings transformed from image patches is fed into multiple Swin Transformer blocks and patch merging layers, with Swin Transformer blocks performing feature learning, and patch merging layers downsampling the feature resolution and unifying the feature dimension. The designed bottleneck comprises two consecutive Swin Transformer blocks to learn the hierarchical representation from the encoder with feature resolution and dimension unchanged.

Swin Transformer blocks and patch-expanding layers construct the symmetric Transformer-based decoder. In contrast to the patch merging layers in the encoder, each patch expanding layer is responsible for upsampling the feature maps into double resolutions and halving the corresponding feature dimension. The final reshaped feature maps pass through a linear projection to produce the pixel-wise segmentation outputs. Inspired by the U-Net, the framework also employs skip connections to combine multi-scale features with the upsampled features at various resolution levels to reduce the loss of fine-grained contextual information caused by down-sampling.

In contrast to the CNN-based methods showing over-segmentation issues, the proposed U-shape pure Transformer presents better segmentation performance resulting from learning both local and long-range dependencies. Compared to the previous methods [173, 17], the HD evaluation metric of Swin-Unet shows an improvement in accuracy for better edge prediction. The experiments on the Synapse multi-organ CT dataset and ACDC dataset from MRI scanners also demonstrate the robustness and generalization ability of the method.

Compared to Swin-Unet and DS-TransUNet, **nnFormer** [153] proposed by Zhou et al. preserves the superior performance of convolution layers for local detail extraction and employs a hierarchical structure to model multi-scale features. It utilizes the volume-based multi-head self-attention (V-MSA) and the shifted version (SV-MSA) in the Transformer blocks instead of processing 2D slices of the volume. The overall architecture of nnFormer is composed of an encoder and a decoder. Each stage in the encoder and decoder consists of a Transformer block applying V-MSA and SV-MSA and a successive upsampling or downsampling block built upon convolution layers, which is referred to as the interleaved architecture. V-MSA conducts self-attention within 3D local volumes instead of 2D local windows to reduce the computational complexity by approximately 98% and 99.5% on the Synapse and ACDC datasets, respectively. nnFormer is first pre-trained on the ImageNet dataset [73] and utilizes symmetrical initialization to reuse the pre-trained weights of the encoder in the decoder. It employs the self-supervised learning strategy, deep supervision, to train the network. Specifically, the output feature maps of all resolutions in the decoder are passed to the corresponding expanding block. It applies cross-entropy loss and dice loss to evaluate the loss for each stage. Accordingly, the resolution of ground truth is downsampled to match the predictions of all the resolutions. The weighted losses of all three stages comprise the final training objective function. The results of experiments that compare nnFormer with prior Transformer-based [17, 26] and CNN-based arts [174] illustrate nnFormer makes significant progress on the segmentation task.
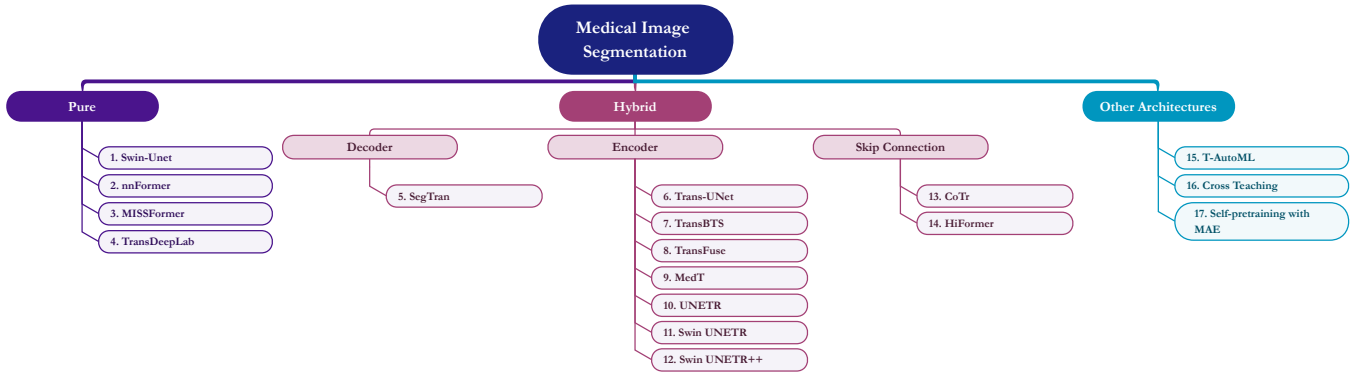
Figure 9: An overview of ViTs in medical image segmentation. Methods are classified into the pure Transformer, hybrid Transformer, and other architectures according to the positions of the Transformers in the entire architecture. The prefix numbers of the methods denote 1. [26], 2. [153], 3. [154], 4. [27], 5. [155], 6. [17], 7. [156], 8. [157], 9. [158], 10. [159], 11. [160], 12. [161], 13. [162], 14. [31], 15. [163], 16. [164], 17. [165].
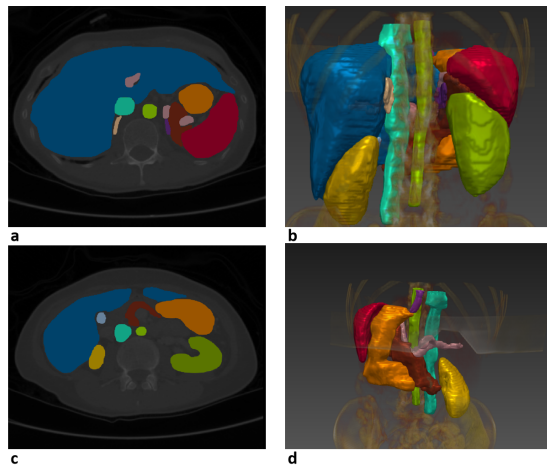
.



Figure 10: Transformer-based models can perform image segmentation on medical image datasets. Figure **a** and **c** illustrate two 2D slices of raw images with the labels from Synapse dataset [170]. Figure **b** and **d** show the 3D visualization of the labeled organs from different angles. These images were generated with MITK Workbench [171].

Although the recent Transformer-based methods improve the problem that CNN methods cannot capture long-range dependencies, they show the limitation of the capability of modeling local details. Some methods directly embedded the convolution layers between fully connected layers in the feed-forward network. Such structure supplements the low-level information but limits the discrimination of features. Huang et al. propose **MISSFormer** [154], a hierarchical encoder-decoder network, which employs the Transformer block named Enhanced Transformer Block and equips the Enhanced Transformer Context Bridge.

The Enhanced Transformer Block utilizes a novel efficient self-attention module that illustrates the effectiveness of spatial reduction for better usage of the high-resolution map. The original multi-head self-attention can be formulated as follows:

$$Attention(Q, K, V) = Softmax(\frac{QK^T}{\sqrt{d_{head}}})V, \quad (4)$$

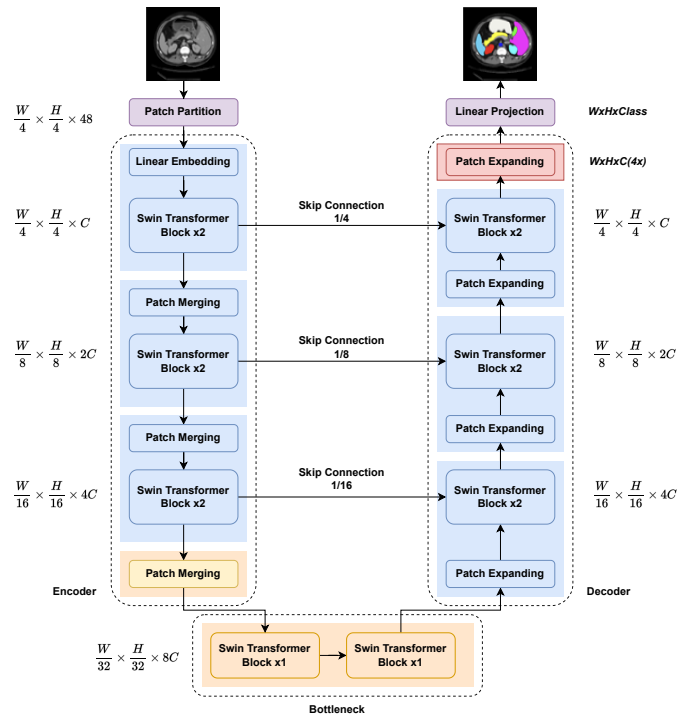where $Q$, $K$, and $V$ refer to query, key, and value, respectively,



Figure 11: The architecture of the Swin-Unet [26] which follows the U-Shape structure. It contains the encoder, the bottleneck, and the decoder part, which are built based on the Swin Transformer block. The encoder and the decoder are connected with skip connections.

and have the same shape of $N \times C$, $d_{head}$ denotes the number of heads. The computational complexity is $O(N^2)$. In efficient self-attention, the $K$ and $V$ are reshaped by a spatial reduction ratio $R$. Take $K$, for example:

$$new\_K = Reshape(\frac{N}{R}, C \cdot R)W(C \cdot R, C) \quad (5)$$

$K$ is first resized from $N \times C$ to $\frac{N}{R} \times (C \cdot R)$ and then projected linearly to restore the channel depth from $C \cdot R$ to $C$. The computational cost reduces to $O(\frac{N^2}{R})$ accordingly. Furthermore, the structure of the Enhanced Mix Feed-forward network (Mix-FFN) extended from [175] introduces recursive skip con-

nections to make the model more expressive and consistent with each recursive step.

The U-shaped architecture of the MISSFormer contains the encoder and decoder built on the Enhanced Transformer blocks connected with an enhanced Transformer context bridge. Multi-scale features produced from the encoder are flattened and concatenated together and passed through the Enhanced Transformer Context Bridge. The pipeline of the Enhanced Transformer Context Bridge is based on the Enhanced Transformer Block to fuse the hierarchical features. The output of the bridge is split and recovered to each original spatial dimension to pass through the corresponding stage in the decoder. The results of experiments show a robust capacity of the method to capture more discriminative details in medical image segmentation. It is worth mentioning that MISSFormer trained from scratch even outperforms state-of-the-art methods pre-trained on ImageNet.

The results in Figure 12 show that the performance of MISSFormer for prediction and segmentation of edges in pure Transformer network structures is more accurate compared to TransUNet and Swin-Unet. Comparing MISSFormer and MISSFormer-S (MISSFormer without bridge), MISSFormer has fewer segmentation errors because the bridge is effective for integrating multi-scale information.



(a) GroundTruth   (b) MISSFormer   (c) MISSFormer_S   (d) Swin-Unet   (e) TransUnet
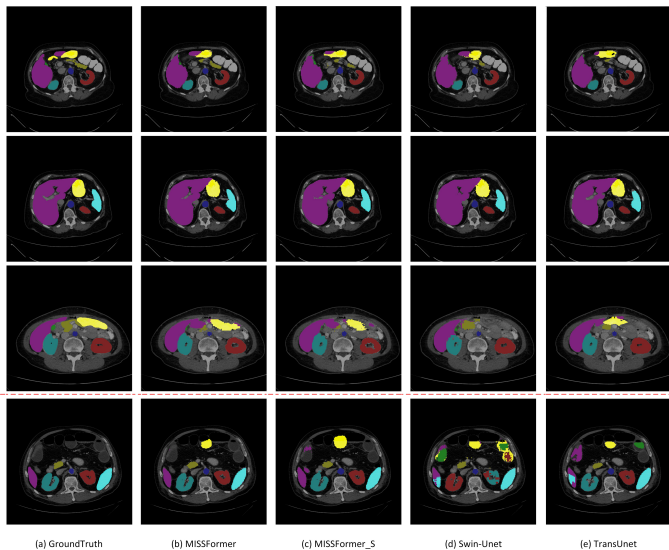
Figure 12: A visual comparison with the state-of-the-art approaches on Synapse dataset. Above the red line shows the successful cases of segmentation, and below the red line are the failed cases with relatively large errors [154]

.

Inspired by the notable DeepLabv3 [176] which utilizes the Atrous Spatial Pyramid Pooling (ASPP) to learn multi-scale feature representations and depth-wise separable convolution to reduce the computational burden, the authors propose **TransDeepLab** [27] to combine the DeepLab network with the Swin-Transformer blocks. Applying the Swin-Transformer module with windows of multiple sizes enables the fusion of multi-scale information with a lightweight model.

TransDeepLab is a pure Transformer-based DeepLabv3+ architecture, as shown in Figure 13. The model builds a hierarchical architecture based on the Swin-Transformer mod-

ules. TransDeepLab first employs $N$ stacked Swin-Transformer blocks to model the input embedded images into a deep-level feature space. 2D medical images are first to split into non-overlapping patches of dimension $C$ and size $4 \times 4$. The ensuing Swin-Transformer block learns local semantic information and global contextual dependencies of the sequence of patches. Then, the authors introduce windows with different sizes to process the output of the Swin-Transformer and fuse the resulting multi-scale feature layers, which are then passed through Cross Contextual attention. This design, referred to as Spatial Pyramid Pooling (SSPP) block, replaces the original Atrous Spatial Pyramid Pooling (ASPP) module exploited in DeepLabV3. A cross-contextual attention mechanism is utilized to explore the multi-scale representation after fusion. This attention module applies channel attention and spatial attention to the output from windows of each size (from each layer of the spatial pyramid). Finally, in the decoder part, the low-level features from the encoder are concatenated with the multi-scale features extracted by Cross Contextual Attention after bilinear upsampling. The last two Swin-Transformer blocks and the patch expanding module generate the final prediction masks.
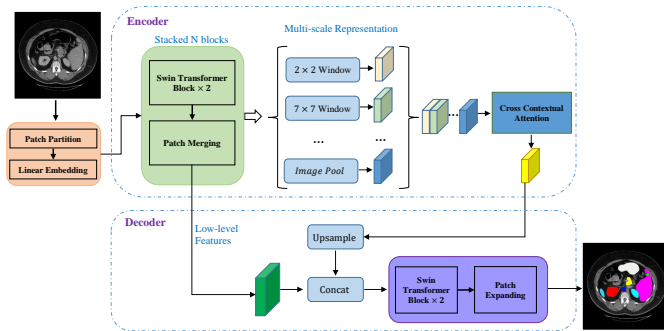


Figure 13: The overview architecture of TransDeepLab, which comprises encoder and decoder built on Swin-Transformer blocks. It is the pure Transformer-based extension of DeepLabv3++ [27].

### 4.2. Hybrid Models

Hybrid Transformers concatenate Transformer blocks with convolution layers to extract local details and long-range dependencies. We further classify this category into *Transformer: Encoder*, *Transformer: Decoder* and *Transformer: skip connection* according to the position of the combined module in the U-Net architecture.

### 4.2.1. Transformer: Encoder

Starting with TransUNet [17], multiple methods in the medical image segmentation field adopt the self-attention mechanism in the encoder.

Transformers have developed as an alternative architecture for modeling global context that exclusively relies on attention mechanisms instead of convolution operators. However, its inner global self-attention mechanism induces missing low-level details. Direct upsampling cannot retrieve the local information, which results in inaccurate segmentation results. The authors propose the **TransUNet** architecture, a hybrid approach that integrates CNN-Transformer hybrid as the encoder and

cascaded upsampler as the decoder, combining the advantages of Transformer and U-Net to boost the segmentation performance by recovering localized spatial information.

The framework of the TransUNet is illustrated in Figure 14. The proposed encoder initially employs CNN as a feature extractor to build a feature map for the Transformer input layer, rather than the Transformer directly projecting the raw tokenized picture patches to latent embedding space. In this way, the intermediate CNN feature maps of different resolutions can be saved and utilized in the following process.

For the decoder, the Cascaded Upsampler (CUP) is proposed to replace naive bilinear upsampling, applying several upsampling blocks to decode the hidden feature and output the final segmentation result. Finally, the hybrid encoder and the CUP constitute the overall architecture with skip connections to facilitate feature aggregation at different resolution levels. This strategy can compensate for the loss of local fine-grained details caused by the Transformer encoder and merge the encoded global information with the local information contained in intermediate CNN feature maps.

The experiments show that TransUNet significantly outperforms the model consisting of pure Transformer encoder and naive upsampling, as well as the ViT-hybrid model without skip connections [17]. Comparisons with prior work [28, 177] also demonstrate the superiority of TransUNet over competing CNN-based approaches in terms of both qualitative visualization and the quantitative evaluation criteria (i.e.average DSC and HD). TransUNet integrates the benefits of both high-level global contextual information and low-level details as an alternative approach for medical image segmentation.
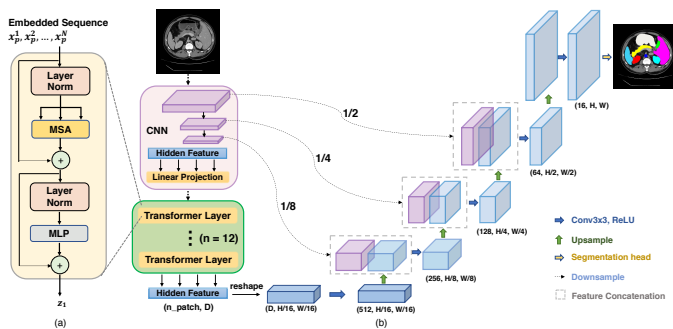


Figure 14: The overview architecture of the TransUNet [17]. The Transformer layers are employed in the encoder part. The schematic of the Transformer is shown on the left.

Wang et al. [156] propose the encoder-decoder architecture, **TransBTS**, which leverages Transformer on learning global contextual information and merits the 3D CNN for modeling local details. In contrast to the concurrent Transformer-based model [17], which analyzes 3D medical volumetric data in a slice-by-slice manner, TransBTS also explores the local features along the depth dimension by processing all the image slices at once.

The network encoder initially employs a 3D CNN to capture volumetric spatial features, simultaneously downsampling the input 3D images, yielding compact volumetric feature maps. Each feature map is projected into a token and fed into the

Transformer encoder to investigate the global relationships. The full-resolution segmentation maps are generated by the 3D CNN decoder after the progressive upsampling while using the feature embedding from the Transformer. For the encoder part, TransBTS first utilizes the $3 \times 3 \times 3$ convolution blocks with downsampling to process the 3D input medical image data, which boosts the effective embedding of rich local 3D features a cross spatial and depth dimensions into the low-resolution feature representation $F$. They apply a linear projection to the feature representation $F$ to obtain the sequence $f$, which is then integrated with position embeddings, as the input for the Transformer encoder. The Transformer encoder consists of multiple Transformer layers, each of which comprises a Multi-Head Attention(MHA) block and a Feed-Forward Network(FFN). The output sequence of the Transformer encoder passes through the feature mapping module to be reshaped to a 4D feature map $Z$ of the same dimension as $F$. The approach employs cascaded upsampling and convolution blocks to progressively restore the segmentation predictions at the original resolution. Furthermore, skip connections combine the fine-grained details of local information with the decoder modules, resulting in more accurate segmentation masks.

The authors conduct comparisons between the proposed TransBTS and the closest method TransUNet [17]. TransUNet essentially processes 3D medical images slice by slice, while TransBTS is a 3D model that explores the continuous interaction through the depth dimension by processing a 3D medical image in a single pass. In contrast to TransUNet, which adopts pre-trained ViT models on other large-scale datasets, TransBTS is trained on the dataset for the specified task without relying on pre-trained weights.

The framework is evaluated on the Brain Tumor Segmentation (BraTS) 2019 challenge and 2020 challenge. Compared to the 3D U-Net baseline, TransBTS achieves a significant enhancement in segmentation. The prediction results indicate the improved accuracy and the superiority of modeling long-range dependencies.

Previous approaches [157] primarily focus on replacing convolution operation with Transformer layers or consecutively stacking the two together to address the inherent lack of pure Transformer-based models to learn local information. In this study, the authors propose a new strategy-**TransFuse** which consists of the CNN-based encoder branch and the Transformer-based branch in parallel fused with the proposed BiFusion module, thus further exploring the benefits of CNNs and Transformers. The construction of the Transformer branch is designed in the typical encoder-decoder manner. The input images are first split into non-overlapped patches. The linear embedding layer then projects the flattened patches into the raw embedding sequence which is added to the learnable position embedding of the same dimension. The obtained embeddings are fed into the Transformer encoder, which comprises $L$ layers of MSA and MLP. The output of the last layer of the Transformer encoder is passed through layer normalization to obtain the encoded sequence.

The decoder part utilizes the same progressive upsampling (PUP) approach as SETR [178]. The encoded sequence is first

reshaped back to a sequence of 2D feature maps. Then they employ two stacked upsampling-convolution layers to restore the feature scales. The feature maps with different spatial resolutions generated by each upsampling-convolution layer are retained for the subsequent fusion operation. For the CNN branch, the approach discards the last layer of the traditional CNNs architecture and combines the information extracted from the CNNs with the global contextual features obtained from the Transformer branch. A shallower model is yielded as a result of this design, avoiding the requirement for extremely deep models that exhaust resources to get long-range dependencies. For instance, there are five blocks in a typical ResNet-based network where only the outputs of the 4th, 3rd, and 2nd layers are saved for the following fusion with the feature maps from the Transformer branch.

The BiFusion module is proposed to fuse the features extracted from the two branches mentioned above to predict the segmentation results of medical images. The global features from the Transformer branch are boosted by the channel attention proposed in SE-Block [9]. Meanwhile, the feature maps from the CNN branch are filtered by the spatial attention which is adopted in the CBAM [179] block to suppress the irrelevant and noisy part and highlight local interaction. Then the Hadamard product is applied to the features from the two branches to learn the interaction between them. They concatenate the interaction feature $b^i$ with attended features $\hat{t}^i$ and $\hat{g}^i$ and feed the results through a residual block to produce the feature $f^i$, which successfully models both the global and local features at the original resolution. Finally, the segmentation prediction is generated by integrating the $f^i$ from different BiFusion modules via the attention-gated (AG) [180] skip connection.

They evaluate the performance of three variants of Trans-Fuse on four segmentation tasks with different imaging modalities and target sizes. TransFuse-S is constructed with ResNet-34 (R34) and 8-layer DeiT-Small (DeiT-S) [65]. Besides, TransFuse-L is composed of Res2Net-50 and 10-layer DeiT-Base (DeiT-B). TransFuse-L* is implemented based on ResNetV2-50 and ViT-B [15]. For polyp segmentation, Transfuse-S/L outperforms significantly the CNN baseline models with fewer parameters and faster running time. TransFuse-L* also achieves the best performance among the previous SOTA Transformer-based methods with a faster speed for inference. It runs at 45.3 FPS and about 12% faster than TransUNet. The experiments for other segmentation tasks also show the superiority of the segmentation performance.

Despite the powerful results of applying Transformers to segmentation tasks [181, 178], the dilemma is that properly training existing Transformer-based models requires large-scale datasets, whereas the number of images and labels available for medical image segmentation is relatively limited. To overcome the difficulty, **MedT** [158] proposes a gated position-sensitive axial attention mechanism where the introduced gates are learnable parameters to enable the model to be applied to a dataset of arbitrary size. Furthermore, they suggested a Local-Global(LoGo) training strategy to improve the segmentation performance by operating on both the original image and the local patches.
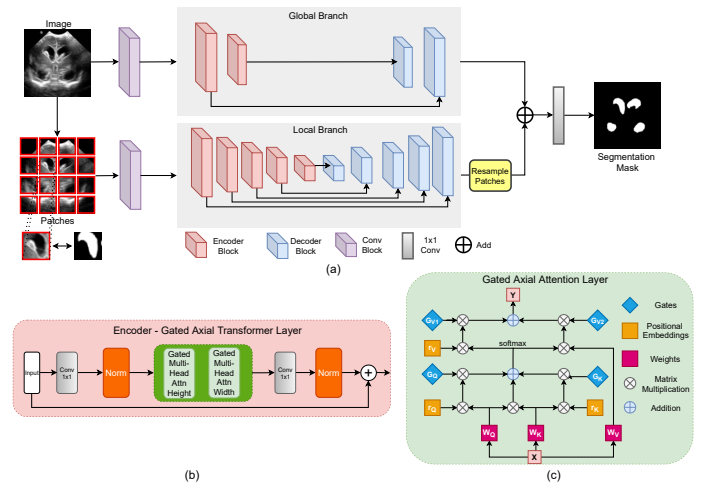


Figure 15: Overview of the MedT [158] architecture. The network uses the LoGo strategy for training. The upper global branch utilizes the first fewer blocks of the Transformer layers to encode the long-range dependency of the original image. In the local branch, the images are converted into small patches and then fed into the network to model the local details within each patch. The output of the local branch is re-sampled relying on the location information. Finally, a $1 \times 1$ convolution layer fuses the output feature maps from the two branches to generate the final segmentation mask.

The main architecture of MedT, as shown in Figure 15 (a), is composed of 2 branches: a shallow global branch that works on the original resolution of the entire image, and a deep local branch that acts on the image patches. Two encoder blocks and two decoder blocks comprise the global branch, which is sufficient to model long-range dependencies. In the local branch, the original image is partitioned into 16 patches and each patch is feed-forwarded through the network. The output feature maps are re-sampled based on their locations to obtain the output feature maps of the branch. Then the results generated from both branches are added and fed into a $1 \times 1$ convolution layer to produce the output segmentation mask. The LoGo training strategy enables the global branch to concentrate on high-level information and allows the local branch to learn the finer interactions between pixels within the patch, resulting in improved segmentation performance.

Figure 15 (b) and (c) illustrate the gated axial Transformer layer, which is used as the main building block in MedT, and the feed-forward structure in it. They introduced four learnable gates $G_{V1}, G_{V2}, G_Q, G_K \in \mathbb{R}$ that control the amount of information the positional embeddings supply to key, query, and value. Based on whether a relative positional encoding is learned accurately or not, the gate parameters will be assigned weights either converging to 1 or some lower value. The gated mechanism can control the impact of relative positional encodings on the encoding of non-local context and allows the model to work well on any dataset regardless of size.

Unlike the fully-attended baseline [181], MedT trained on even smaller datasets outperforms the convolutional baseline and other Transformer-based methods. In addition, improvements in medical segmentation are also observed since the proposed method takes into account pixel-level dependencies.

In contrast to multiple proposed methods [178, 17, 158, 157]

that investigate the task of 2D medical image segmentation, **UNETR** [159] proposes a novel Transformer-based architecture for 3D segmentation which employs the Transformer as the encoder to learn global contextual information from the volumetric data. In addition, unlike the previous frameworks proposed for 3D medical image segmentation [162, 156], the encoded feature from the Transformer of this proposed model is directly connected to a CNN-based decoder via skip connections at different resolution levels. The U-shaped UNETR comprises a stack of Transformers as the encoder and a decoder coupling with it by skip connections. They begin by generating the 1D sequence of patches by splitting the 3D input volume in a non-overlapping manner. The flattened input patches are then passed through a linear projection layer to yield $K$ dimensional patch embeddings. They attach a 1D learnable positional embedding to each patch embedding taking into account the spatial information of the extracted patches. After the embedding layer, the global multi-scale representation is captured using Transformer blocks composed of multi-head self-attention modules and multilayer perceptron layers. They resize and project the sequence representation extracted from the Transformer at different resolutions for use in the decoder in order to retrieve spatial information of the low-level details.

In the expanding pattern of the framework, the proposed CNN-based decoder combines the output feature of different resolutions from the Transformer with upsampled feature maps to properly predict the voxel-wise segmentation mask at the original input resolution.

The paper claims UNETR achieves new state-of-the-art performance on all organs compared against CNN [182, 29, 183, 184] and competing for Transformer-based [162, 17, 178] baselines on BTCV dataset, with significant improvement in performance on small organs in particular. In addition, it outperforms the closest methodologies on brain tumor and spleen segmentation tasks in the MSD dataset. UNETR shows the superiority of learning both global dependencies and fine-grained local relationships in medical images.

Figure 16 presents qualitative segmentation comparisons for brain tumor segmentation on the MSD dataset between UNETR [159], TransBTS [156], CoTr [162] and U-Net [28]. It can be seen that the details of the brain tumor are captured well by UNETR [159].

As opposed to other methods that attempted to utilize the Transformer module as an additional block beside the CNN-based components in the architectures, UNETR [159] leverages the Transformer as the encoder instead of the CNN-based encoder. The Swin Transformer [58] is a hierarchical visual Transformer featuring an efficient shift-window partitioning scheme for computing self-attention. Inspired by these two approaches, a novel model termed **Swin Unet Tr**ansformer (**Swin UNETR**) [160] is proposed for brain tumor segmentation in this work.

The proposed framework applies a U-shape architecture with the Swin Transformers as the encoder and a CNN-based module as the decoder connected to the encoder via skip connections at different resolution levels. The model initially converts 3D MRI images with four channels to non-overlapping patches
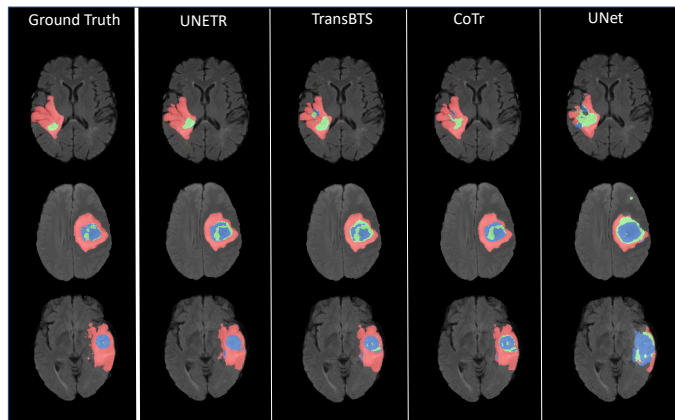


Figure 16: Comparison of visualization of brain tumor segmentation on the MSD dataset. The whole tumor (WT) includes a combination of red, blue, and green regions. The union of red and blue regions demonstrates the tumor core (TC). The green regions indicate the enhanced tumor core (ET) [159].

and creates windows of a specific size with a patch partition layer.

The Swin UNETR encoder is composed of 4 stages. Each stage comprises 2 Transformer blocks and a patch merging layer. In the Transformer blocks, the self-attention is computed with a shifted windowing mechanism. Swin UNETR employs the windows of size $M \times M \times M$ to partition the 3D token with resolution of $H' \times W' \times D'$ into regions of $\lceil \frac{H'}{M} \times \frac{W'}{M} \times \frac{D'}{M} \rceil$ at layer $l$. The partitioned window regions are then shifted by $(\lfloor \frac{M}{2} \rfloor, \lfloor \frac{M}{2} \rfloor, \lfloor \frac{M}{2} \rfloor)$ voxels at the following $l + 1$ layer. The patch merging layer after the Transformer components reduces the resolution of feature maps by a factor of two and concatenates them to form a feature embedding with the doubled dimensionality of the embedding space.

For the decoder of the architecture, the output feature representations of the bottleneck are reshaped and passed through the residual block containing two convolutional layers. The subsequent deconvolutional layer increases the resolution of feature maps by a factor of 2. The outputs are then concatenated with the outputs of the previous stage and fed into another residual block. After the resolutions of the feature maps are restored to the original $H' \times W' \times D'$, a head is utilized to generate the final segmentation predictions.

The authors conduct the experiments to compare Swin UNETR against the previous methodologies SegResNet [185], nn-UNet [174] and TransBTS [156] in this work. The results demonstrate that the proposed model has prominence as one of the top-ranking approaches in the BraTS 2021 challenge. It is due to the better capability of learning multi-scale contextual information and modeling long-range dependencies by Swin Transformers in comparison to regular Transformers with a fixed resolution of windows.

Apart from the architecture of Swin UNETR [160], they also proposed a hierarchical encoder for pre-training the encoder of Swin UNETR [161], called **Swin UNETR++**. The pre-training tasks are tailored to learning the underlying pattern of human anatomy. Various proxy tasks such as image inpainting, 3D rotation prediction, and contrastive learning are

leveraged to extract the features from the naturally consistent contextual information in 3D radiographic images such as CT. Specifically, inpainting enables learning the texture, structure, and relationship of masked regions to their surrounding context. Contrastive learning is utilized to recognize separate regions of interest (ROIs) of different body compositions. The rotation task creates a variety of sub-volumes that can be used for contrastive learning. To validate the effectiveness of the pre-trained encoder, they fine-tuned the pre-trained Swin UNETR on the downstream segmentation task and achieved SOTA performance on BTCV dataset [170].

### 4.2.2. Transformer: Decoder

Another direction is to modify the decoder of the U-shape structure to aggregate the Transformer-CNN-based modules.

In the **Segtran** framework [155], as illustrated in Figure 17, Squeeze-and-Expansion Transformer is proposed to "squeeze" the attention matrix and aggregate multiple sets of contextualized features from the output. A novel Learnable Sinusoidal Position Encoding is also employed to impose the continuity inductive bias for images. The Segtran consists of five components: a CNN backbone to extract image features, 2) input/output feature pyramids to do upsampling, 3) the Learnable Sinusoidal Positional Encoding, 4) Squeeze-and-Expansion Transformer layers to contextualize features, and 5) a segmentation head. The pretrained CNN backbone is first utilized to learn feature maps from the input medical images. Since the input features to Transformers are of a low spatial resolution, the authors increase their spatial resolutions with an input Feature Pyramid Network (FPN) [186] to upsample the feature maps by bilinear interpolation. Then the proposed Learnable Sinusoidal Positional Encoding is added to the visual features to inject spatial information. In contrast to the previous two mainstream PE schemes [187, 15], the new positional embedding vector, a combination of sine and cosine functions of linear transformations of $(x, y)$, brings in the continuity bias with adaptability. The equation of the encoding strategy varies gradually with pixel coordinates. Thus, close units receive similar positional encodings, increasing the attention weights between them towards higher values. The encoding vectors generated from the addition of positional encodings and visual features are then fed into the Transformer.

The novel Transformer architecture combines Squeezed Attention Block (SAB) [188] with an Expanded Attention Block. Here this method employs the Induced Set Attention Block (ISAB) proposed by [188] as a squeezed attention block. The Squeezed Attention Block computes attention between the input and inducing points and compresses the attention matrices to lower rank matrices, reducing noises and overfitting. The Expanded Attention Block (EAB), a mixture-of-experts model, outputs $N_m$ sets of complete contextualized features from $N_m$ modes. Each mode is an individual single-head Transformer and shares the same feature space with each other. That is as opposed to multi-head attention in which each head outputs an exclusive feature subset. All features are then aggregated into one set using dynamic mode attention. The dynamic mode attention can be obtained by doing a linear transformation of each mode feature and taking softmax over all the modes.

Compared with representative existing methods in the experiments, Segtran consistently achieved the highest segmentation accuracy and exhibited good cross-domain generalization capabilities.
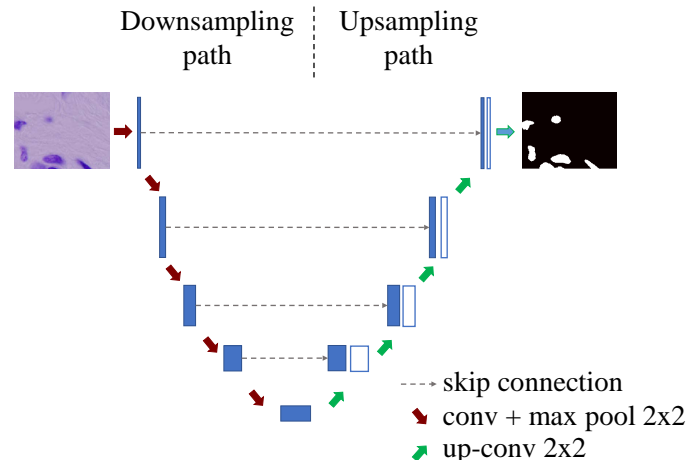


Figure 17: Segtran network extracts image features using a CNN backbone and combines the features with the position encoding of pixels flattened into a series of local feature vectors. Multiple squeezed and extended transform layers are stacked to process the local feature vectors. Finally, an output FPN after the Transformer upsamples the features to generate the final prediction [155].

### 4.2.3. Transformer: Skip Connection

In this section, Transformer blocks are incorporated into the skip connections to facilitate the transmission of detailed information from the encoder to the decoder.

Although Transformer-based methods overcome the limitation of capturing long-range dependency, they present extreme computational and spatial complexity in analyzing high-resolution volumetric image data. Some studies [187, 17] employ hybrid structures, fusing CNN with Transformer in an attempt to reduce the training requirement on huge datasets. The recent approach, TransUNet [17], shows good performance. However, it is difficult to optimize the model due to the inner self-attention mechanism of the vanilla Transformer. First, it takes a long time to train the attention, which is caused by initially distributing attention uniformly to each pixel within the salient regions [14]. Second, a vanilla Transformer struggles to handle multi-scale and high-resolution feature maps due to its high computational cost.

Motivated by this, [162] proposes a novel encoder-decoder framework, **CoTr**, which bridges CNN and Transformer. The architecture exploits CNN to learn feature representations. An efficient deformable self-attention mechanism in the Transformer is designed to model the global context from the extracted feature maps, which reduces the computational complexity and enables the model to process high-resolution features. The final segmentation results are generated by the decoder.

As shown in Figure 18, the DeTrans-encoder consists of an input-to-sequence layer and multiple DeTrans Layers. The input-to-sequence layer first flattens the feature maps at different resolutions extracted from the CNN-encoder into 1D se-

quences $\{f_l\}_{l=1}^L$. Then the corresponding 3D positional encoding sequence $p_l$ is supplemented with each of the flattened sequences $f_l$ to complement the spatial information. The combined sequence is fed as the input into the DeTrans Layers. Each of the DeTrans Layers is a composition of an MS-DMSA and a Feed-Forward Network (FFN). In contrast to the self-attention mechanism which casts attention to all the possible locations, the proposed MS-DMSA layer only attends to a small set of key sampling locations around a reference location. As a result, it can achieve faster convergence and lower computational complexity. The skip connection is utilized after each DeTrans Layer to preserve the low-level details of local information. The output of the DeTrans-encoder is successively upsampled by the pure CNN encoder to restore the original resolution. Besides, they apply skip connections and a deep supervision strategy to add fine-grained details and auxiliary losses to the prediction outputs.

The results of experiments indicate that CoTr with the hybrid architecture has superior of performance over the models with pure CNN encoder or pure Transformer encoder. It also outperforms other hybrid methods like TransUNet [17] in processing multi-scale 3D medical images with reduced parameters and complexity.
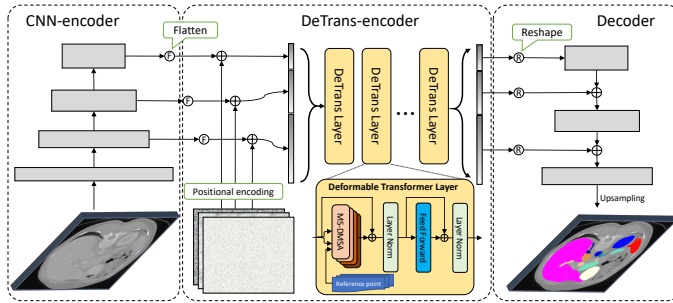


Figure 18: Overview of the CoTr [162] architecture. It is composed of a CNN-encoder, a DeTrans-encoder and a decoder. The CNN-encoder models the local information of the input images and provides the outputs at each stage. The outputs of different resolutions are flattened, fused and passed through the Deformable Transformer Layers along with positional encoding. The decoder reshapes the processed sequences from the DeTrans-encoder and produces the final predictions.

**HiFormer** [31] is proposed to aggregate a fusion module in the skip connections to learn richer representations. Figure 19 demonstrates the end-to-end network structure of the strategy that incorporates the global dependencies learned with the Swin Transformer and the detailed local features extracted by the CNN modules. The encoder is composed of two hierarchical CNN, Swin Transformer modules and the novel Double-Level Fusion module (DLF module). First, medical images are fed into a CNN module to obtain a local fine-grained semantic representation. After the CNN layer catches the shallow feature layers, HiFormer introduces the Swin Transformer modules to complement the global feature information. The Swin Transformer module employs windows of different sizes to learn the dependencies between multiple scales. To reuse the shallow and deep multi-scale feature information in the encoder, HiFormer designs a novel skip connection module, the DLF module. The deep-level semantic and shallow-level lo-

calization information are fed into the DLF module and fused by the cross-attention mechanism. Finally, both generated feature maps are passed into the decoder to produce the final segmentation prediction results. The experiments conducted on the Synapse dataset [170], SegPC [189], and ISIC 2017 dataset [117] demonstrate the superiority of the learning ability of HiFormer. Moreover, the lightweight model with fewer parameters also exceeds CNN-based methods and previous Transformer-based approaches with lower computational complexity.
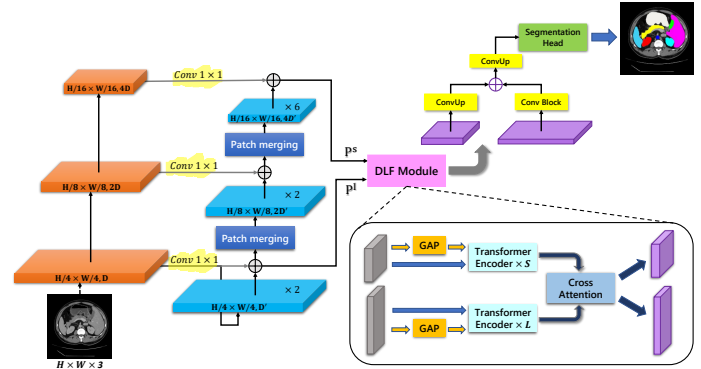


Figure 19: HiFormer comprises the CNN-Transformer encoder, the CNN-based decoder and the Double-Level Fusion Module (DLF). The feature layers of the shallowest level $p^l$ and of the deepest level $p^s$ are fed into the DLF module for the fusion of hierarchical information. Blue blocks and orange blocks refer to Swin Transformer and CNN modules, respectively [31].

### 4.3. Other Architectures

Most ViT-based models rely on pre-training of large natural image datasets to obtain pre-weights and then solve downstream tasks by transfer learning. Several works explore training in a self-supervised or semi-supervised manner to efficiently utilize medical image datasets of limited size or datasets without manual labels. Furthermore, some approaches apply Transformers to seek the design of architectures that implement medical image segmentation, instead of using the Transformers to act directly on the input image.

Unlike the previously proposed methods that employ the Transformers to act directly on the medical image for feature extraction, this method [163] adopts the AutoML for automatically designing the network architecture without much human heuristics or assumptions, where the Transformer is applied to encode the embedding vector regarding the architecture configurations. The approach reduces the workload of algorithm design by automatically estimating "almost" all the components of the framework instead of manually designing for the network and training strategies. That improves the model performance of segmentation simultaneously.

The proposed Transformer-based **T-AutoML** inspired by SpineNet [190] leverages neural architecture search (NAS) with a larger search space to optimize the selection of the network connections. This framework can connect the feature maps at different spatial levels of the network with another one arbitrarily, compared with the previous methods that only search for the encoder-decoder U-shape networks [191, 192, 193]. The candidates of different blocks in the network consist of 3D residual

21

blocks, 3D bottleneck blocks, and 3D axial-attention blocks. The residual blocks and bottleneck blocks are effective in alleviating the vanishing gradient. The axial-attention blocks are applied to model the long-range dependency in the 2D medical images. Another upsampling layer (linear interpolation) is utilized at the end of the architecture to produce the results of feature maps at the original volume size.

To search for the optimal architecture and training configuration, the authors first encode the necessary components in the search space to form a one-dimensional vector $v$. The search space contains candidates of different configurations with regard to data augmentation, learning rates, learning rate schedulers, loss function, the optimizer, the number and spatial resolution of blocks, and block types.

After the obtainment of the encoding vector $v$, the proposed new predictor predicts the binary relation of validation accuracy values between $vi$ and $v_j$. The predictor employs the Transformer encoder to encode the vector $v$ of varying lengths into feature maps of a fixed resolution. Then the feature maps are passed through the Multiple FC layers to generate the binary relation predictions denoted as $GT_{v_i,v_j}$. Since the predictor is designed for ranking the vectors with respect to the accuracy values and estimating the relations, the actual values of the predicted accuracy are not necessary to be calculated for each vector. Thus, the new predictor requires less overall training time.

The experiments indicate that the proposed method can achieve the state of the art(SOTA) in lesion segmentation tasks and shows superiority in transferring to different datasets.

Despite the promising results achieved by the CNNs and Transformer-based methods with large-scale images, these approaches require expert labeling at the pixel/voxel level. Expensive time costs and manual annotation limit the size of the medical image dataset. Due to this dilemma, the proposed semi-supervised segmentation [164] provides a low-cost and practical scheme, called **Cross Teaching** between CNN and Transformer, to train effective models using a little amount of correctly labeled data and a large amount of unlabeled or coarsely labeled data.

Inspired by the existing works [194, 195, 196] for semi-supervised learning which introduce perturbation at different levels and encourage prediction to be consistent during the training stage, the designed cross teaching introduces the perturbation in both learning paradigm-level and output-level. As shown in Figure 20, each image within the training set containing labeled and unlabeled images is fed into two different learning paradigms: a CNN and a Transformer respectively. For the unlabeled dataset with raw images, the cross teaching scheme allows the cross supervision between a CNN ($f_\phi^c(.)$) and a Transformer($f_\phi^t(.)$), which aims at integrating the properties of the Transformer modeling the long-range dependency and CNN t0 learn local information in the output level.

The unlabeled data initially passes through a CNN and a Transformer respectively to generate predictions $p_i^c$ and $p_i^t$.

$$p_i^c = f_\phi^c(x_i); p_i^t = f_\phi^t(x_i); \qquad (6)$$

Then the pseudo labels $pl_i^c$ and $pl_i^t$ are produced in this manner:

$$pl_i^c = argmax(p_i^t); pl_i^t = argmax(p_i^c); \qquad (7)$$

The pseudo label $pl_i^c$ used for the CNN training is generated by the Transformer. Similarly, the CNN model provides pseudo labels for Transformer training. The cross-teaching loss for the unlabeled data is defined as follows:

$$L_{ctl} = \underbrace{L_{Dice}(p_i^c, pl_i^c)}_{supervision\ for\ CNNs} + \underbrace{L_{Dice}(p_i^t, pl_i^t)}_{supervision\ for\ Transformers} \qquad (8)$$

which is a bidirectional loss function. One direction of the data stream is from the CNN to the Transformer, and the other direction is from the Transformer to the CNN. For the labeled data, the CNN and Transformer are supervised by the ground truth. The commonly-used supervised loss functions, i.e. the cross-entropy loss and Dice loss, are employed to update model parameters.

$$L_{sup} = L_{ce}(p_i, y_i) + L_{Dice}(p_i, y_i) \qquad (9)$$

where $p_i$ , $y_i$ represent the prediction and the label of image $x_i$. And the overall objective combining the cross-teaching branch and supervised branch is defined as:

$$L_{total} = L_{sup} + \lambda L_{ctl} \qquad (10)$$

where $\lambda$ is a weight factor, which is defined by time-dependent Gaussian warming up function [197, 164]:

$$\lambda(t) = 0.1 \cdot e^{\left(-5\left(1-\frac{t_i}{t_{total}}\right)\right)^2} \qquad (11)$$

The results of the ablation study indicate that the combination of CNN and Transformer in a cross-teaching way shows superiority over the existing semi-supervised methods. Furthermore, the novel method has the potential to reduce the label cost by learning from limited data and large-scale unlabeled data. However, it is observed that achieving SOTA via semi-supervised approaches remains a significant challenge.
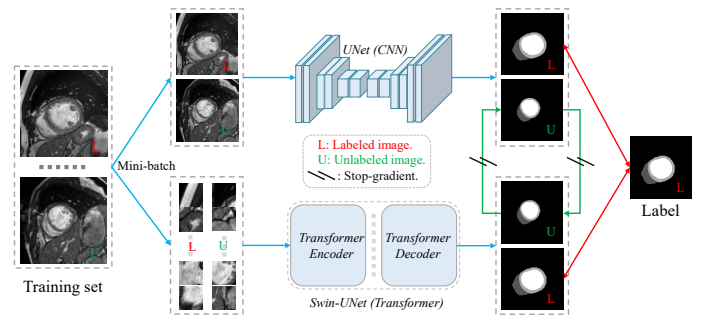


Figure 20: The model performs the semi-supervised medical image segmentation task. The regularization scheme between CNN and Transformer is referred to as Cross Teaching. L denotes the labeled data and U denotes the unlabeled data. The cross-teaching employs a bidirectional loss function: one path is from the CNN branch to the Transformer branch, and the other is from the Transformer to the CNN. A Transformer is applied for complementary training instead of prediction generation [164].

Zhou et al. [165] hypothesize that the ability to aggregate contextual information is imperative to improve the performance of medical image analysis. Nonetheless, there is no ImageNet-scale medical image dataset for pre-training. Therefore, they investigate a novel self-pre-training paradigm based

Table 4: An overview of the reviewed Transformer-based medical image Segmentation models.

| Method | Modality | Organ | Type | Pre-trained Module: Type | Datasets | Metrics | Year |
|---|---|---|---|---|---|---|---|
| **Pure** | | | | | | | |
| Swin-Unet [26] | CT | Multi-organ | 2D | ViT: Supervised | [1] Synapse [170] [2] ACDC [198] | Dice | 2021 |
| nnFormer [153] | CT MRI | Multi-organ | 3D | ViT: Supervised | [1] Synapse [170] [2] ACDC [198] | Dice | 2021 |
| MISSFormer [154] | CT MRI | Multi-organ | 2D | ✗ | [1] Synapse [170] [2] ACDC [198] | Dice Hausdorff distance | 2021 |
| TransDeepLab [27] | CT Dermoscopy | Multi-organ Skin | 2D | ViT: Supervised | [1] Synapse [170] [2] ISIC 2017, 2018 [117, 118] [3] PH2 [199] | Dice Hausdorff distance | 2022 |
| **Encoder** | | | | | | | |
| TransUNet [17] | CT MRI | Multi-organ | 2D | ViT: Supervised | [1] Synapse [170] [2] ACDC [198] | Dice Hausdorff distance | 2021 |
| TransBTS [156] | MRI | Brain | 3D | ✗ | BraTS 19-20 [200, 201, 202] | Dice Hausdorff distance | 2021 |
| TransFuse [157] | Colonoscopy | Multi-organ | 2D 3D | ViT: Supervised | [1] Synapse [170] [2] ACDC [198] | Dice Hausdorff distance | 2021 |
| MedT [158] | Microscopy Ultrasound | Multi-organ | 2D | ✗ | [1] Brain US (Private) 2 GLAS [203] [3] MoNuSeg [204] | F1 | 2021 |
| UNETR [159] | CT MRI | Brain, Spleen | 3D | ✗ | [1] Synapse [170] [2] MSD [205] | Dice Hausdorff distance | 2021 |
| Swin UNETR [160] | MRI | Brain | 3D | ViT: Supervised | BraTS 21 [206] | Dice Hausdorff distance | 2022 |
| Swin UNETR++ [161] | CT | Multi-organ | 3D | ViT: Self-supervised | [1] BTCV [170] [2] MSD [205] | Dice NSD | 2022 |
| **Skip connection** | | | | | | | |
| CoTr [162] | CT | Multi-organ | 3D | ✗ | Synapse [170] | Dice | 2021 |
| HiFormer [31] | MRI Dermoscopy Microscopic | Multi-organ Skin Cells | 2D | ViT: Supervised CNN: Supervised | [1] Synapse [170] [2] ISIC 2017, 2018 [117, 118] [3] SegPC 2021 [207, 208, 209] | Dice Hausdorff distance | 2022 |
| **Decoder** | | | | | | | |
| SegTran [155] | Fundus MRI X-Colonoscopy | Multi-organ | 2D 3D | CNN: Supervised | REFUGE 20 [210] [1] BraTS 19 [200, 201, 202] [2] X-CVC [211] [3] KVASIR [212] | Dice | 2021 |
| **Other architectures** | | | | | | | |
| T-AutoML [163] | CT | Liver and lung tumor | 3D | ✗ | MSD 2019 [213] | Dice NSD | 2021 |
| Cross Teaching [164] | MRI | Multi-organ | 2D | ✗ | ACDC [198] | Dice Hausdorff distance | 2022 |
| Self-pretraining with MAE [165] | CT MRI X-ray | Lung Brain Multi-organ | 3D | ViT: supervised | [1] ChestX-ray14 [104] [2] Synapse [213] [3] MSD 2019 [205] | Dice Hausdorff distance | 2022 |

on Masked Autoencoder (MAE), **MAE self pre-training**, for medical image analysis, one of the masked image modeling (MIM) frameworks [215] [216] [217] [218]. MIM encourages the framework to restore the masked target by integrating information from the context, where the main idea of MIM is masking and reconstructing: masking a set of image patches before input into the Transformer and reconstructing these masked patches at the output.

The pipeline for segmentation with MAE self-pre-training contains two stages. In the first stage (as shown on the left of Figure 21), ViT is pre-trained with MAE as the encoder. The input patches are divided into visible ones and masked ones. The ViT encoder only acts on the visible patches. Compared to other MIM methods, MAE does not employ mask tokens in the encoder, which saves time and allows for faster pre-training. A lightweight Transformer decoder is appended to reconstruct the full image. The decoder is only an auxiliary part used for pre-training and will not be applied in downstream tasks.

In the second stage (as shown on the right of Figure 21), the

pre-trained ViT weights are transferred to initialize the segmentation encoder. Then, the task-specific heads are appended to perform downstream tasks. The whole segmentation network, e.g., UNETR, is finetuned to perform the segmentation task.

The experiments, including the MAE pre-training and the downstream task, are conducted to evaluate the performance of the proposed method. The results show that MAE can recover the lost information in the masked input patches. MAE pre-training enables the model to improve its classification and segmentation performance on medical image analysis tasks, surpassing the ImageNet pre-trained model to SOTA.

### 4.4. Discussion and Conclusion

This section comprehensively investigates the overview of around 17 Transformer-based models for medical image segmentation presented in Section 4.1 to Section 4.3. Some representative self-supervised strategies [161, 153] for train-

Table 5: A brief description of the reviewed Transformer-based medical image segmentation models. The unreported number of parameters indicates that the value was not mentioned in the paper, and the code was unavailable.

| Method | # Params | Contributions | Highlights |
|---|---|---|---|
| **Pure** | | | |
| Swin-Unet[26] | - | • Builds a pure Transformer model with symmetric Encoder-Decoder architecture based on the Swin-Transformer block connected via skip connections.<br>• Proposes patch merging layers and patch expanding layers to perform downsampling and upsampling without convolution or interpolation operation. | • The results of extensive experiments on multi-organ and multi-modal datasets show the good generalization ability of the model.<br>• Pre-trained on ImageNet rather than medical image data, which may result in sub-optimal performance. |
| nnFormer [153] | 158.92M | • Proposes a powerful segmentation model with an interleaved architecture (stem) based on the empirical combination of self-attention and convolution.<br>• Proposes a volume-based multi-head self-attention (V-MSA) to reduce computational complexity. | • Volume-based operations help to reduce the computational complexity.<br>• Pre-trained on ImageNet rather than medical image data, which may result in sub-optimal performance. |
| MISSFormer [154] | - | • Proposes the Enhanced Transformer Block based on the Enhanced Mix-FFN and the Efficient Self-attention module.<br>• Proposes the Enhanced Transformer Context Bridge built on the Enhanced Transformer Block to model both the local and global feature representation and fuse multi-scale features. | • The model can be trained from scratch without the pretraining step on ImageNet.<br>• Less computational burden due to the novel design of the Efficient Self-attention module. |
| TransDeepLab [27] | 21.14M | • Proposes the encoder-decoder DeepLabv3+ architecture based on Swin-Transformer.<br>• Proposes the cross-contextual attention to adaptively fuse multi-scale representation. | • The first attempt to combine the Swin-Transformer with DeepLab architecture for medical image segmentation.<br>• A lightweight model with only 21.14M parameters compared with Swin-Unet[26], the original DeepLab model [214] and TransUNet[17]. |
| **Encoder** | | | |
| TransUNet [17] | 96.07M | • Proposes the first CNN-Transformer hybrid network for medical image segmentation, which establishes self-attention mechanisms from the perspective of sequence-to-sequence prediction.<br>• Proposes a cascaded upsampler (CUP) which comprises several upsampling blocks to generate the prediction results. | • TransUNet fully exploits the strong global context encoded from the Transformer and local semantics from the CNN module.<br>• It presents the generalization ability on multi-modalities.<br>• The approach allows the segmentation of 2D and 3D medical images. |
| TransBTS [156] | Moderate TransBTS: 32.99M<br>Lightweight TransBTS: 15.14M | • Proposes a novel encoder-decoder framework TransBTS that integrates Transformer with 3D CNN for MRI Brain Tumor Segmentation. | • The method can model the long-range dependencies not only in spatial but also in the depth dimension for 3D volumetric segmentation.<br>• TransBTS can be trained on the task-specific dataset without the dependence on pre-trained weights.<br>• TransBTS is a moderate-size model that outperforms in terms of model complexity with 32.99M parameters and 33G FLOPs. Furthermore, the vanilla Transformer can be replaced with other variants to reduce the computation complexity. |
| TransFuse [157] | TransFuse-S: 26.3M | • Proposes the first parallel-in-branch architecture — TransFuse to capture both low-level global features and high-level fine-grained semantic details.<br>• Proposes BiFusion module in order to fuse the feature representation from the Transformer branch with the CNN branch. | • The architecture does not require very deep nets, which alleviates gradient vanishing and feature diminishing reuse problems.<br>• It improves performance by reducing parameters and increasing inference speed, allowing deployment on both the cloud and the edge.<br>• The CNN branch is flexible to use any off-the-shelf CNN network.<br>• It can be applied to both 2D and 3D medical image segmentation. |
| MedT [158] | 1.4M | • Proposes a gated axial-attention model that introduces an additional control mechanism to the self-attention module.<br>• Proposes a LoGo (Local-Global) training strategy for boosting segmentation performance by simultaneously training a shallow global branch and a deep local branch. | • The proposed method does not require pre-training on large-scale datasets compared to other transform-based models.<br>• The results of predictions are more precise compared to the full attention model. |
| UNETR [159] | 92.58M | • Proposes a novel architecture to address the task of 3D volumetric medical image segmentation.<br>• Proposes a new architecture where the Transformer-based encoder learns long-range dependencies and the CNN-based decoder utilizes skip connections to merge the outputs of a Transformer at each resolution level with the upsampling part. | • UNETR shows moderate model complexity while outperforming these Transformer-based and CNN-based models.<br>• The inference time of UNETR is significantly faster than Transformer-based models.<br>• They did not use any pre-trained weights for the Transformer backbone. |
| Swin UNETR [160] | 61.98M | • Proposes a novel segmentation model, Swin UNETR, based on the design of UNETR and Swin Transformers. | • The FLOPs of Swin UNETR significantly grow compared to that of UNETR and TransBTS.<br>• The Swin Transformer is suitable for the downstream tasks. |
| Swin UNETR++ [160] | 61.98M | • Presents a new self-supervised learning approach that capitalizes on the Swin UNETR model, harnessing its multi-resolution encoder for efficient pre-training.<br>• Introduces a pre-trained robust model using the proposed encoder and proxy tasks on a dataset of 5,050 publicly available CT images, resulting in a powerful feature representation for various medical image analysis tasks. | • Fine-tuning the pretrained Swin UNETR model leads to improved accuracy, faster convergence, and reduced annotation effort compared to training from randomly initialized weights.<br>• Researchers can benefit from the pre-trained encoder in transfer learning for various medical imaging analysis tasks, including classification and detection. |
| **Skip connection** | | | |
| CoTr [162] | 46.51M | • Proposes a hybrid framework that bridges a convolutional neural network and a Transformer for accurate 3D medical image segmentation.<br>• Proposes the deformable Transformer (DeTrans) that employs the multi-scale deformable self-attention mechanism (MS-DMSA) to model the long-range dependency efficiently. | • The deformable mechanism in CoTr reduces computational and spatial complexities, allowing the network to model high-resolution and multi-scale feature maps. |
| HiFormer [31] | HiFormer-S: 23.25M<br>HiFormer-B: 25.51M<br>HiFormer-L: 29.52M | • Proposes an encoder-decoder architecture that bridges a CNN and a Transformer for medical image segmentation.<br>• Proposes a Double-level Fusion module in the skip connection. | • Fewer parameters and lower computational cost. |
| **Decoder** | | | |
| SegTran [155] | 86.03M | • Proposes a novel Transformer design, Squeeze-and-Expansion Transformer, in which a squeezed attention block helps regularize the huge attention matrix, and an expansion block learns diversified representations.<br>• Proposes a learnable sinusoidal positional encoding that imposes a continuity inductive bias for the Transformer. | • Compared to U-Net and DeepLabV3+, Segtran has the least performance degradation, showing the best cross-domain generalization when evaluated on datasets of drastically different characteristics. |
| **Other architectures** | | | |
| T-AutoML [163] | 16.96M | • Proposes the first automated machine learning algorithm, T-AutoML, which automatically estimates "almost" all components of a deep learning solution for lesion segmentation in 3D medical images.<br>• Proposes a new predictor-based search method in a new search space that searches for the best neural architecture and the best combination of hyperparameters and data augmentation strategies simultaneously. | • The method is effectively transferable to different datasets.<br>• The applied AutoML alleviates the need for the manual design of network structures.<br>• The intrinsic limitations of AutoML. |
| Cross Teaching [164] | - | • Proposes an efficient regularization scheme for semi-supervised medical image segmentation where the prediction of a network serves as the pseudo label to supervise the other network.<br>• The proposed method is the first attempt to apply the Transformer to perform the semi-supervised medical segmentation utilizing the unlabeled data. | • The training process requires less data cost by semi-supervision.<br>• The framework contains components with low complexity and simple training strategies compared to other semi-supervised learning methods.<br>• The proposed semi-supervised segmentation still can not achieve the state-of-the-art (SOTA) compared with the fully-supervised approaches. |
| Self-pretraining with MAE [165] | - | • Proposes a self-pre-training paradigm with MAE for medical images where the pre-training process of the model uses the same data as the target dataset. | • The proposed paradigm demonstrates its effectiveness in limited data scenarios. |

Table 6: Comparison results of some segmentation methods on the *Synapse* [170] dataset. The methods are ordered based on their dice score.

| Methods | Params (M) | FLOPs (G) | Dice ↑ | HD95 ↓ | Aorta | Gallbladder | Kidney(L) | Kidney(R) | Liver | Pancreas | Spleen | Stomach |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| **TransUNet [17]** | 96.07 | 88.91 | 77.49 | 31.69 | 87.23 | 63.16 | 81.87 | 77.02 | 94.08 | 55.86 | 85.08 | 75.62 |
| **Swin-Unet [26]** | 27.17 | **6.16** | 79.13 | 21.55 | 85.47 | 66.53 | 83.28 | 79.61 | 94.29 | 56.58 | 90.66 | 76.60 |
| **TransDeepLab [27]** | **21.14** | 16.31 | 80.16 | 21.25 | 86.04 | 69.16 | 84.08 | 79.88 | 93.53 | 61.19 | 89.00 | 78.40 |
| **HiFormer [31]** | 25.51 | 8.05 | 80.39 | 14.70 | 86.21 | 65.23 | 85.23 | 79.77 | 94.61 | 59.52 | 90.99 | 81.08 |
| **PVT-CASCADE [219]** | 35.28 | 6.40 | 81.06 | 20.23 | 83.01 | 70.59 | 82.23 | 80.37 | 94.08 | 64.43 | 90.10 | 83.69 |
| **MISSFormer [154]** | 42.46 | 9.89 | 81.96 | 18.20 | 86.99 | 68.65 | 85.21 | 82.00 | 94.41 | 65.67 | 91.92 | 80.81 |
| **DAEFormer [220]** | 48.07 | 27.89 | 82.63 | 16.39 | 87.84 | 71.65 | 87.66 | 82.39 | 95.08 | 63.93 | 91.82 | 80.77 |
| **TransCASCADE [219]** | 123.49 | - | 82.68 | 17.34 | 86.63 | 68.48 | 87.66 | 84.56 | 94.43 | 65.33 | 90.79 | 83.52 |
| **ScaleFormer [221]** | 111.60 | 48.93 | 82.86 | 16.81 | **88.73** | **74.97** | 86.36 | 83.31 | 95.12 | 64.85 | 89.40 | 80.14 |
| **D-LKA Net [222]** | 101.64 | 19.92 | 84.27 | 20.04 | 88.34 | 73.79 | **88.38** | **84.92** | 94.88 | 67.71 | 91.22 | 84.94 |
| **MERIT [223]** | 147.86 | 33.31 | **84.90** | **13.22** | 87.71 | 74.40 | 87.79 | 84.85 | **95.26** | **71.81** | 92.01 | 85.38 |

Table 7: Performance comparison of several Transformer-based segmentation approaches on *ISIC 2017* [117], *ISIC 2018* [118], and *PH²* [224].

| Methods | Params (M) | FLOPs (G) | ISIC 2017 | | | | ISIC 2018 | | | | PH² | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | DSC | SE | SP | ACC | DSC | SE | SP | ACC | DSC | SE | SP | ACC |
| **TransUNet [17]** | 96.07 | 88.91 | 0.8123 | 0.8263 | 0.9577 | 0.9207 | 0.8499 | 0.8578 | 0.9653 | 0.9452 | 0.8840 | 0.9063 | 0.9427 | 0.9200 |
| **TransNorm [172]** | 117.98 | 33.41 | 0.8933 | 0.8532 | **0.9859** | 0.9582 | 0.8951 | 0.8750 | 0.9790 | 0.9580 | 0.9437 | **0.9438** | **0.9810** | **0.9723** |
| **Swin-Unet [26]** | 27.17 | 6.16 | 0.9183 | 0.9142 | 0.9798 | 0.9701 | 0.8946 | 0.9056 | **0.9798** | 0.9645 | 0.9449 | 0.9410 | 0.9564 | 0.9678 |
| **D-LKA Net [222]** | 101.64 | 19.92 | **0.9254** | **0.9327** | 0.9793 | **0.9705** | **0.9177** | **0.9164** | 0.9773 | **0.9647** | **0.9490** | 0.9430 | 0.9775 | 0.9659 |
| **HiFormer [31]** | 25.51 | 8.05 | 0.9253 | 0.9155 | 0.9840 | 0.9702 | 0.9102 | 0.9119 | 0.9755 | 0.9621 | 0.9460 | 0.9420 | 0.9772 | 0.9661 |
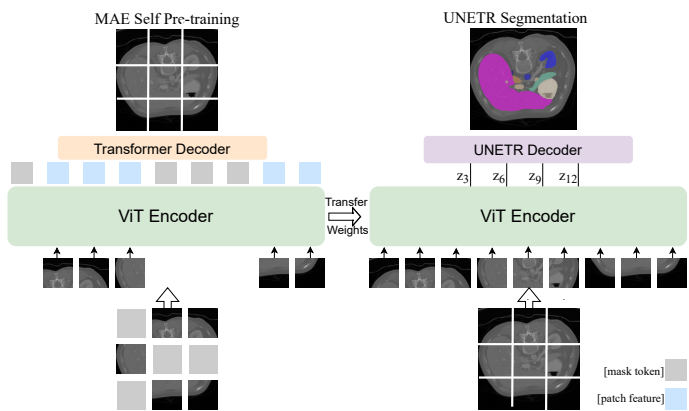
Figure 21: Illustration of MAE self pre-training. First, MAE is pre-trained as an encoder for ViT. The ViT encoder is fed with a random subset of patches and the decoder of the Transformer reconstructs the complete image as shown on the left. Then, the pre-trained ViT weights are transferred to the initialized segmentation encoder, as shown on the right. Finally, the whole segmentation network, such as UNETR, is fine-tuned to perform the segmentation task. [165].

ing transformer modules are also introduced in this section. We provide information on the reviewed segmentation approaches about the architecture type, modality, organ, input size, the pre-trained manner, datasets, metrics, and the year in Table 4. Table 5 also lists the methods along with the number of parameters, contributions, and highlights. The comparison results of some SOTA methods on the Synapse dataset are also presented in Table 6, showcasing their respective performance metrics. Additionally, the evaluation of SOTA approaches on the ISIC 2017, ISIC 2018, and PH$^2$ datasets are demonstrated in Table 7. The tables serve as valuable references for evaluating the relative performance and suitability of different approaches in medical image segmentation. Moreover, inference time and GPU usage are crucial performance factors for model selection. DAE-Former excels with a fast 46.01 ms inference and 6.42 GB GPU memory. MERIT, while effective, takes longer at 93.13 ms and uses 7.78 GB. D-LKA, with 113.18 ms inference, balances GPU memory well at 5.22 GB. TransCASCADE offers a compromise with 53.18 ms inference and 3.75 GB GPU memory. These metrics alongside the parameter numbers and FLOP count guide model selection for efficient medical image segmentation. It's important to note that these assessments were conducted using the same test setup and configurations as those employed in image classification Section 3.2.

ViT-based works offer solutions in a broad range of multimodal tasks of 2D or 3D. Most of the approaches demonstrate superior results over CNN-based segmentation models on benchmark medical datasets. Despite the SOTA performance Transformer-based networks have achieved, there are some challenges in deploying the Transformer-based models at present. The first challenge is the high computational burden due to the relatively large number of parameters of the Transformer-based models [220]. The reason is that the time and space complexity of the attention mechanism is quadratic to the sequence length. For example, the CNN-based models such as U-Net [28] require 3.7M parameters [158] to reach Dice Score 74.68 [17]. However, TransUNet, which achieves Dice Score 77.48 needs 96.07M [159] parameters. The researchers have to meet the high demand for GPU resources. Thus, several novel approaches such as Swin Transformer employed in Swin-Unet [26], volume-based Transformer utilized in nnFormer [153] and efficient self-attention module in MISSFormer [154] are proposed to simplify the computation of the Transformer models. The direction of facilitating the efficiency of models will play a crucial role in future research. We also note that most existing methods require pre-training strategies on the ImageNet dataset to obtain the pre-trained weights for the following downstream tasks. However, the natural image datasets and medical datasets differ dramatically from one another, which may impact the final performance of extracting the medical features. Meanwhile, pre-training leads to high computational costs, which hinders the training of models in practice. Multiple segmentation networks that can be trained from scratch on the medical dataset are suggested as the solutions, such as MISSFormer [154]. We expect more approaches to exploring more efficient pre-training strategies or without pre-training. Furthermore, considering the limited size of some medical datasets, some approaches propose semi-supervised technologies or self-pre-training paradigms to reduce the dataset burden of training or pre-training. Nevertheless, the performance is still not comparable to that of fully-supervised models. Designing semi-supervised models with improved accuracy in this direction requires more attention.

## 5. Medical Image Reconstruction

3D Medical imaging is a clinical breakthrough and is very popular in medical diagnosis and follow-up after treatment. In Computed Tomography (CT), Single Photon Emission Tomography (SPECT), and Positron Emission Tomography (PET), the imaging process relies on ionizing radiation [242, 243], which implies a potential risk for the patient [244]. A non-invasive 3D imaging technique is Magnetic Resonance Imaging (MRI), which does not rely on ionizing radiation. However, image acquisition may take longer and confines the patient in a discomforting narrow tube [245]. In order to reconstruct 3D volumetric datasets from the acquired data, Medical image reconstruction is one of the essential components of 3D medical imaging. The primary objective of 3D image reconstruction is to generate high-quality volumetric images for clinical usage at minimal cost and radiation exposure, whilst also addressing potential artifacts inherent to the physical acquisition process. Image reconstruction solves an inverse problem that is generally challenging due to its large-scale and ill-posed nature [246].

In medical imaging, there are ongoing research efforts to reduce the acquisition time (i.e. to reduce cost and potential
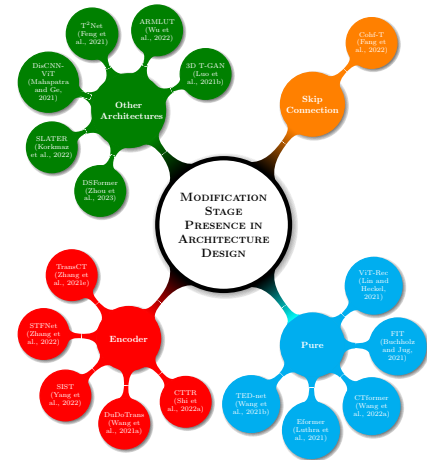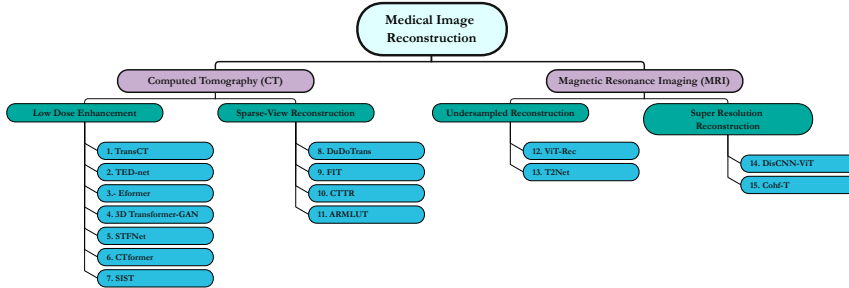
(a) Taxonomy structure for medical image reconstruction. Methods in this field are categorized by their functionality in addressing issues in consensus imaging modalities, not how the Transformer is integrated with the architecture like in the previous sections. The prefix numbers in the paper's name in ascending order denote the reference for each study as follows: 1. [225], 2. [226], 3. [227], 4. [228], 5. [229], 6. [230], 7. [231], 8. [232], 9. [233], 10. [234], 11. [235], 12. [236], 13. [237], 14. [238], 15. [239], 16. [240], 17. [241].

(b) We also presented a second taxonomy due to the presence of the Transformer in the reviewed studies.

Figure 22: An overview of medical image reconstruction taxonomies either as categorizing by the task or the location of using Transformer in an architecture.

movement artifacts) as well as radiation dose. However, lowering the radiation dose results in higher noise levels and reduced contrast, which poses a challenge for 3D image reconstruction.

Vision Transformers (ViTs) have effectively demonstrated possible solutions to address these challenges. We categorize the literature in this domain into *low dose enhancement*, *sparse-view reconstruction*, *undersampled reconstruction*, and *super-resolution reconstruction*. This section will overview some of the SOTA Transformer-based studies that fit into our taxonomy. Figure 22a and Figure 22b demonstrate our proposed taxonomy for this field of study. Figure 22a indicates the diversity of our taxonomy based on the medical imaging modalities we studied in this research. Figure 22b endorses the usage of the Transformer within the overviewed studies' pipelines.

### 5.1. Low Dose Enhancement

Zhang et al. [225] used a very general intuition about image denoising: the noisy image constructed with high-frequency and low-frequency counterparts as $X = X_H + X_L$ in a study, namely, **TransCT**. Zhang et al. [225] claim that the noisy image's low counterpart contains two sub-components of main image content and weakened image textures, which are entirely noise-free. They applied a Gaussian filter on the input image to decompose an image into a high-frequency sub-band and a low-frequency sub-band. After this, they extracted $X_{L_c}$ content features and $X_{L_t}$ latent texture features by applying two shallow CNNs on the low-frequency counterpart of the input image. Simultaneously, they applied a sub-pixel layer on a high-frequency counterpart to transform it into a low-resolution image and extracted embedding features ($X_{H_f}$) by applying a shallow CNN. Then the resultant latent texture features ($X_{L_t}$) and corresponding high-frequency representation are fed to the Transformer for noise removal from a high-frequency representation. Ultimately, they reconstruct the high-quality image piecewise. They showed that the latent texture features are beneficial in screening noise from the high-frequency domain.
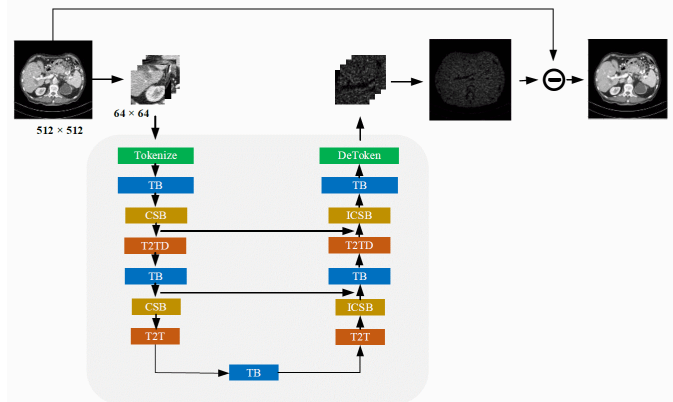


Figure 23: An overview of TED-net [226]. Tokenize and DeToken blocks are invertible operations that apply the process of patch embedding and converting patches again to image, respectively. TB represents a standard Transformer block. (I)CSB denotes the (inverse) Cyclic Shift block to modify the feature map, nonetheless, the reverse operation avoids pixel shifts in the final result. T2T block represents the Token-to-Token process [247] to improve the spatial inductive bias of Transformers by merging the neighboring tokens. The Dilated T2T (T2TD) block is used to refine contextual information further.

Despite the TransCT [225], Wang et al. proposed a convolution-free Token-to-Token vision Transformer-based Encoder-decoder Dilation network (**TED-net**) design for CT image denoising [226]. Their approach is based on a U-Net encoder-decoder scheme enriched by different modules, i.e., Basic Vision Transformer, Token-to-Token Dilation (T2TD), and (Inverse) Cyclic Shift blocks. Consider $y \in \mathbb{R}^{N \times N}$ a clean natural dose CT image, $x \in \mathbb{R}^{N \times N}$ noisy low dose image, and $T : \mathbb{R}^{N \times N} \to \mathbb{R}^{N \times N}$ is a Transformer-based denoising model. According to the Figure 23 after tokenization of $x$ and passing through the Vision Transformer block to capture long-range dependencies and alleviate the absence of local inductive bias in Transformers, they employed Token-to-Token serialization [247]. Also, they utilized feature re-assembling

26

with a Cyclic Shift block (CSB) to integrate more information. Obvious from Figure 23, all of these blocks are replicated in a symmetric decoder path, but instead of the CSB, the Inverse Cyclic Shift block (ICSB) is implemented to avoid pixel shifts in the final denoising results ($y = x + T(x)$). They reached SOTA results compared to CNN-based methods and a competitive benchmark with regard to the TransCT [225].

Luthra et al. [227] proposed a Transformer-based network, **Eformer**, to deal with low-dose CT images while concurrently using the edge enhancement paradigm to deliver more accurate and realistic denoised representations. Their architecture builds upon the LeWin (Locally-enhanced Window) Transformer block [248], which is accompanied by an edge enhancement module. The success of the Swin Transformer [58] in capturing the long-range dependencies with the window-based self-attention technique makes it a cornerstone in designing new Transformer blocks due to its linear computational complexity. LeWin Transformer is one of these blocks that capture the global contextual information and, due to the presence of a depth-wise block in its structure, could also capture a local context. Eformer's first step is through the Sobel edge enhancement filter. In every encoder-decoder stage, convolutional features pass through the LeWin Transformer block, and downsampling and upsampling procedures are done by convolution and deconvolution layers. Eformer's learning paradigm is a residual learning scheme, meaning it learns the noise representation rather than a denoised image due to the ease of optimization in predicting a residual mapping.

Akin to Low Dose CT (LDCT), Low-Dose PET (LDPET) is preferable to avoid the radiation risk, especially for cancer patients with a weakened immune system who require multiple PET scans during their treatment at the cost of sacrificing diagnosis accuracy in Standard-Dose PET (SDPET). Luo et al. [228] proposed an end-to-end Generative Adversarial Network (GAN) based method integrated with a Transformer block, namely **3D Transformer-GAN**, to reconstruct SDPET images from the corresponding LDPET images. To alleviate the inter-slice discontinuity problem of existing 2D methods, they designed their network to work with 3D PET data. Analogous to any GAN network, they used a generator network, encoder-decoder, with a Transformer placed in the bottleneck of the generator network to capture contextual information. Due to the computational overhead of Transformers, they did not build their proposed method solely on it. Therefore, they were satisfied to place a Transformer counterpart across CNN layers of the generator to guarantee to extract low-level spatial feature extraction and global semantic dependencies. They also introduced adversarial loss term to their voxel-wise estimation error to produce more realistic images.

In contradiction with other works, Zhang et al. [229] proposed leveraging the PET/MRI data simultaneously for denoising low-count PET images, which is a crucial assessment for cancer treatment. PET scan is an emission Computed Tomography (CT) operating by positron annihilation radiation. Due to the foundation and requirements of PET scans, there is a severe risk of getting infected with secondary cancer by radiotracers. So to degrade the side effects of this imaging process, there

are two potential methods: reduction in radiotracer dose and lessening the patient's bedtime duration. The aforementioned approaches, without a doubt, affect the imaging result quality with decreased contrast to noise ratio and bias in texture. The traditional low-count PET denoising approaches are based on Non-Local Means (NLM) [249], Block Matching 3D (BM3D) [250], and Iterative methods [251], etc., which are firmly in bond with hyperparameter tuning for new data or result in unnatural smoothings over denoised images. Zhang et al. [229] testify that simultaneous PET/MRI could boost one modality in terms of correct attenuation, motion, and partial volume effects, and also, due to the high contrast among soft tissues in MRI, the denoising process of PET images is preferably straightforward. **STFNet** [229] is a U-Net based structure with different medications. They proposed a new Siamese encoder comprising dual input flow for each modality in the encoding path. To obtain sufficient features from different modalities, they used the Spatial Adaptive (SA) block, a dual path in each block with the residual block design, which consists of different consecutive convolutional blocks and deformable convolution with fusion modulation. This module aims to learn more contextual features from each modality. To leverage global attention, they used a Transformer to produce a pixel-to-pixel interaction between the PET and the MRI modality. After this integration, the fused features are input to the two branches based on residual convolution blocks for PET denoising.

Wang et al. [230] proposed the enhancement for their previous work, TED-net [226] convolution-free, solely Transformer-based network, namely **CTformer**. From Figure 24, it is apparent that their network is an unsupervised residual learning, U-Net-like encoder-decoder structure, rather than direct map learning of LDCT to Natural Dose CT (NDCT). The CTformer tries to compensate for the Transformers' deficiency in capturing path inter-boundary information and spatial inductive bias with token rearrangement, T2T [247]. To do so, analogously like TED-net, they used dilation and cyclic shift blocks in the Token2Token block to broaden the receptive field to capture more contextual information and not increase the computational cost.

Table 8: Comparison result on NIH-AAPM-Mayo [252] dataset in low dose enhancement task. *LDE* indicates the Low Dose Enhancement task.

| Methods | Task | Dataset | SSIM ↑ | RMSE ↓ |
|---|---|---|---|---|
| **Eformer** [227] | LDE | NIH-AAPM-Mayo [252] | **0.9861** | **0.0067** |
| **TransCT** [225] | LDE | NIH-AAPM-Mayo [252] | 0.923 | 22.123 |
| **CTformer** [230] | LDE | NIH-AAPM-Mayo [252] | 0.9121 | 9.0233 |

Yang et al. [231] were inspired by how sinogram works and proposed Singoram Inner-Structure Transformer (**SIST**) (Figure 25). This inner structure of the sinogram contains the unique characteristics of the sinogram domain. To do so, they mimic the global and local characteristics of sinogram in a loss function based on sinogram inner-structure, namely Sino-
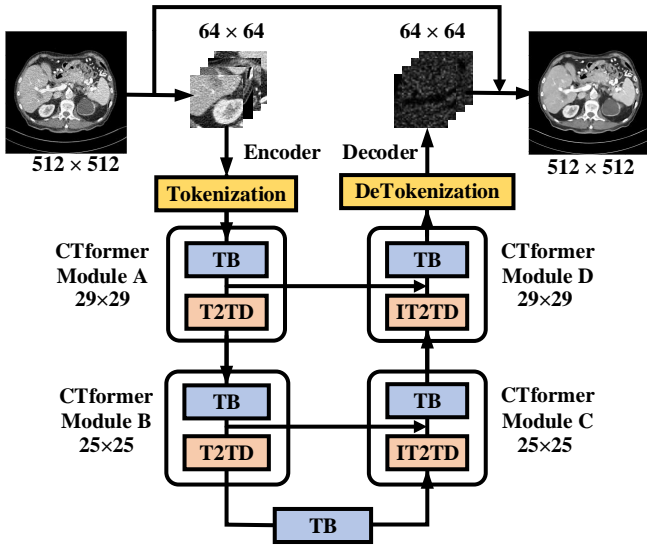
Figure 24: An overview of CTformer [230]. This structure is analogous to the TED-net [226] structure, the previous study by the same authors.

gram Inner-Structure Loss (SISL). The global inner-structure loss utilizes conjugate sampling pairs in CT, and local inner-structure loss considers the second-order sparsity of sinograms. The amalgamation of these two terms could be beneficial in reconstructing NDCT images while retaining the noise. Due to the CT imaging mechanism, each row of the sinogram representation denotes the projection at a certain view. Naturally, this procedure is suitable for leveraging the Transformer block for modeling the interaction between different projections of diverse angles to capture contextual information. Therefore, the SIST module applies to raw sinogram input and captures structural information. Afterward, the unified network reconstructs the high-quality images in a residual policy with the image reconstruction module.
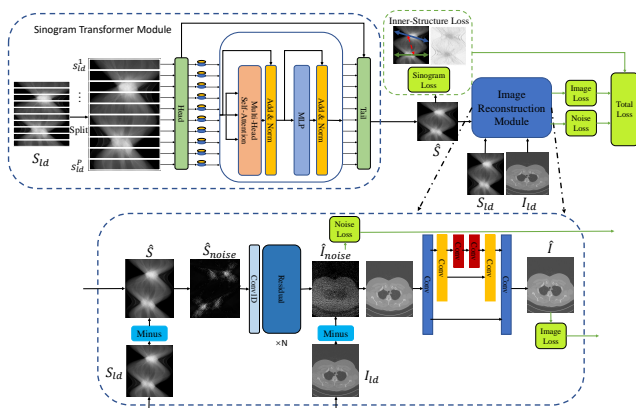


Figure 25: The overall architecture of SIST [231] pipeline. $S_{ld}$ and $I_{ld}$ are the LDCT sinogram and image, $\hat{S}$ and $\hat{I}$ denote the output sinogram and image, $\hat{S}_{noise}$ and $\hat{I}_{noise}$ are the sinogram noise and image noise. First, the LDCT sinogram feed to the Transformer for sinogram domain denoising, then the denoised sinogram $\hat{S}$ input to the image reconstruction module for image domain denoising. Within the image reconstruction module, the sinogram noise $\hat{S}_{noise}$ with the usage of residual CNN block generates image domain $\hat{I}_{noise}$. NDCT $\hat{I}$, outputs from applying refinement steps on $I_{ld}$ minus $\hat{I}_{noise}$.

Table 8 represents the benchmark results in the LDCT task over the NIH-AAPM-Mayo [252] dataset respecting SSIM and RMSE metrics on overviewed methods in this study. For clarification, TED-net [226] achieved better results than CTformer [230], but due to two studies originating from the same authors and the resemblance between architectures, we preferred to mention CTformer to count in the comparison table. This result endorses the capability of the pure Transformer-based Eformer [227] method in reconstructing natural dose CT images.

### 5.2. Sparse-View Reconstruction

Due to the customary usage of CT images in medical diagnosis, another policy to lessen the side effects of X-ray radiation is acquiring fewer projections, known as sparse-view CT, which is a very feasible and effective method rather than manipulating the standard radiation dose [253, 254]. However, the resultant images from this method suffer from severe artifacts, and decreasing the number of projections demands profound techniques to reconstruct high-quality images. Wang et al. [232] is the first paper that inspected the usage of Transformers in this field which was quite successful, namely **DuDo-Trans**. Their intuition was to shed light on the globality nature of the sinogram sampling process, which the previous CNN architectures neglected. DuDoTrans, unlike the conventional iterative methods in this literature, does not provide blocky effects in reconstructed images. This method simultaneously benefits from enhanced and raw sinogram streams to restore informative sinograms via long-range dependency modeling in a supervised policy. DuDoTrans from Figure 26 is built on three main modules, namely Singoram Restoration Transformer (SRT), the DuDo Consistency layer, and the Residual Image Reconstruction Module (RIRM). SRT block consists of successive hybrid Swin Transformer modules and convolutional layers to model local semantic features and inherent global contextual information in the sinogram to produce the enhanced sinogram.
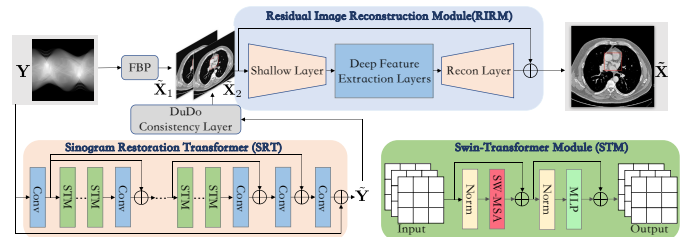


Figure 26: DuDoTrans [232] framework for sparse-view CT image reconstruction. First, the sparse-view sinogram $\mathbf{Y}$ maps to a low-quality image $\bar{\mathbf{X}}_1$ and other estimation $\bar{\mathbf{X}}_2$ generated by SRT module's enhanced sinogram output $\bar{\mathbf{Y}}$ followed by DuDo Consistency Layer. Lastly, the predicted estimations are concatenated and fed to the RIRM module that outputs the CT image of $\bar{\mathbf{X}}$ in a supervised manner.

Buchholz et al. [233] presented the Fourier Image Transformer (**FIT**) that operates on the image frequency representation, especially the Fourier description of the image, which in their study is known as Fourier Domain Encoding (FDE), that encodes the entire image at lower resolution. The intuition in their idea is underlying the CT's acquisition process physics. CT utilizes a rotating 1D detector array around the patient body

to calculate the Radon transform [255] of a 2D object, which leads to a sequence of density measurements at different projection angles, namely sinogram as a 2D image in which each column of this representation corresponds to one 1D measurement. The Filtered Back Projection (FBP) [256, 255] is a reconstruction method to map sinograms to tangible CT images. FBP is based on the Fourier slice theorem; hence, computing the 1D Fourier transform of 1D projection and rearranging them by their projection angle in Fourier space, followed by an inverse Fourier transformation, results in a reconstructed 2D CT image slice. Limiting the number of projections leads to missing Fourier measurements, which ultimately conduce to reconstruction artifacts. FIT is the first study that uses a Transformer to query arbitrary Fourier coefficients and fill the unobserved Fourier coefficients to conceal or avoid the probable artifacts in reconstruction within sparse-view CT reconstruction literature. From Figure 27 this procedure starts with calculating the FDE of the raw sinogram. To do so, first, the discrete Fourier transform (DFT) of the sinogram will be calculated. Secondly, after dropping half of the coefficients on the Fourier rings of the resultant Fourier representation, it preserves the lower frequency counterparts to recover the lower resolution of the raw sinogram. Afterward, the complex coefficients convert into 1D sequences by unrolling the Fourier rings. These complex values convert to normalized amplitudes and phases. Therefore, each complex coefficient has its own polar representation, which is a normalized real-valued matrix with $N \times 2$ entries ($N$ is equal to half of the DFT coefficients number). A linear layer applies on this tensor to upsample the feature dimensionality to $\frac{F}{2}$. Finally, a 2D positional encoding concatenates to this tensor and produces a 2D FDE image with the size of $N \times F$.
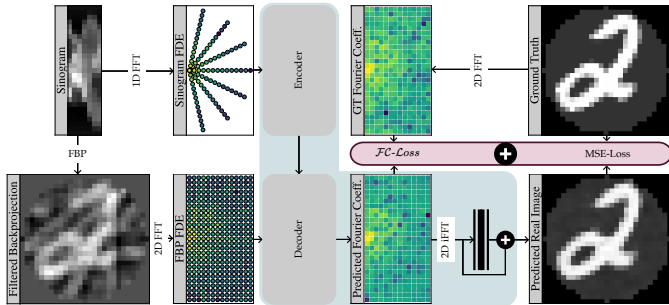


Figure 27: FIT [233] framework for sparse-view CT reconstruction. FDE representation of the sinogram calculates that serves as an input to an encoder of Transformer design. The decoder predicts the Fourier coefficients from the encoder's latent space. The Fourier coefficients of applying the FBP [256] algorithm on sinogram information are fed into a Transformer's decoder to enrich the Fourier query representation. A shallow CNN block applies after inverse FFT to hamper the frequency oscillations.

Shi et al. [234] presented a CT reconstruction network with Transformers (**CTTR**) for sparse-view CT reconstruction. In contrast to DuDoTrans [232], CTTR enhances low-quality reconstructions directly from raw sinograms and focuses on global features in a simple policy in an end-to-end architecture. CTTR contains four parts: two CNN-based residual blocks extracting local features from FBP [256] images reconstruction and sinograms, an encoder-decoder Transformer for long-range

modeling dependencies, and contextual information between features, and a CNN block to map features to a high-quality reconstruction.

Cone-Beam Computed Tomography (CBCT) is a conventional way of dental and maxillofacial imaging; due to its fast 3D imaging qualifications, its popularity has extended to lung imaging. However, studies approved that its radiation dose is higher than plain radiographs [258] hence sparse-view CBCT could be a suitable method to lower radiation dose. Wu et al. [235] proposed a novel untrained 3D Transformer-based architecture, namely **ARMLUT**, with a multi-level loss function for CBCT reconstruction. While the Transformer module, especially the UNETR [159] in this study, captures long-range contextual information and enhances the resulting image. The intuition behind this strategy is Deep Image Prior (DIP) [259] to succeed in the reconstruction field. From Figure 28a, ARMLUT is an iterative optimization problem between the Image Reconstructor module and Image Generator module to fit a CBCT inverse solver without a large number of data or ground truth images. The multi-level loss function comprises a Mean Squared Error (MSE) and Perceptual Loss (PL) [260] to reconstruct smooth and streak artifact-free outputs. The entire framework (Figure 28) has three main counterparts: Image Reconstructor, Image Generator, and Feature Extractor. Image Reconstructor uses Feldkamp-Davis-Kress (FDK) algorithm [257] to produce a low enhanced reconstruction from $M$-view measurements, and the Image generator module maps the noisy voxel inputs to reconstruct a regularised image. The Feature Extractor module applies the VGG-11 pre-trained network on two representations and produces a PL paradigm. To minimize the distance between these two reconstructions, ARMLUT utilizes an adaptively re-weight multi-loss technique to stabilize the convergence of the Transformer in the optimization.

### 5.3. Undersampled Reconstruction

Magnetic Resonance Imaging (MRI) is a dominant technique for assistive diagnosis. However, due to the physics behind its operation, the scanning time can take longer and be very tedious, affecting the patient experience and leading to inevitable artifacts in images [269]. Hence, reducing the number of MRI measurements can result in faster scan times and artifacts reduction due to the patient's movement at the cost of aliasing artifacts in the image [245].

Lin et al. [236] proposed a comprehensive analytical study to investigate the usage of ViT in a pure (CNN-free modules) and most straightforward Transformer design. This study is evidence of the prominent effect of ViTs in medical image reconstruction. For this work, they adopted the original ViT [15] for image reconstruction by discarding the classification token and replacing the classification head with the reconstruction head, which is comprised of successive Norm and Linear layers to map the Transformer output to a visual image. They performed a complete ablation study with different ViT settings, from the number of stacked Transformers to embedding dimension and the number of heads in Multi-Head Self-Attention (MHSA). Their results were quite effective and proved that trained ViT on sufficient data from natural images like ImageNet or medical
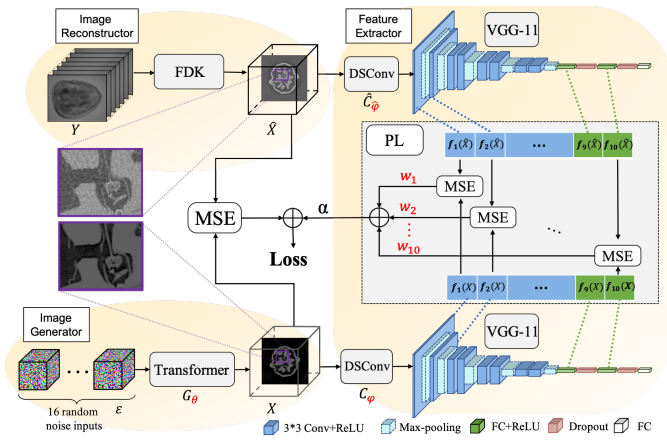
(a) ARMLUT [235] pipeline. The collaboration of three distinct modules—FDK algorithm [257], prior embedding with Transformer, and VGG-11 network for extracting hierarchical features—in this pipeline generates the reconstructed CBCT image. Red texts in the figure denote the variable weights that contribute to the iterative optimization step.
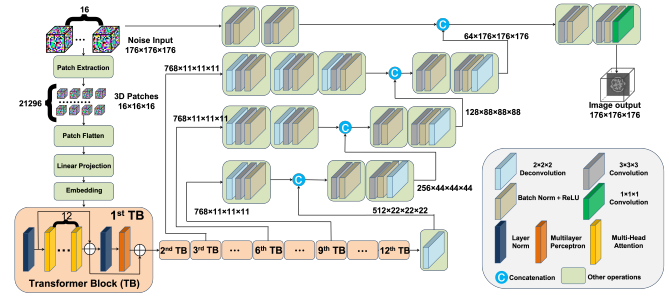


(b) UNETR [159] used as an image generator module in the ARMLUT paradigm.

Figure 28: (a) represents multi-loss untrained network for sparse-view CBCT reconstruction. (b) architecture of UNETR [159], as a Transformer module in a ARMLUT.

data could perform better or achieve on-par reconstruction accuracies compared to CNN baselines such as U-Net [28]. The proposed design's distinguished power based on the mean attention distance metric [270] proves that it effectively mimics the convolutional receptive fields and could concurrently capture local and global dependencies. In addition, they showed that the ViT benefits from two times faster inference times and fewer memory requirements compared to the U-Net.

Feng et al. [237] address the particular issue in this domain by designing an end-to-end multi-task learning paradigm to boost feature learning between two sub-tasks, MRI reconstruction, and super-resolution, which have a high overlap with each other named Task Transformer Network ($T^2Net$) is showed in Figure 29. Their network consists of two branches, each for a specific task. $T^2$Net utilizes a Transformer between two branches to share feature representation and transmission. $T^2$Net applies a convolution layer and EDSR [271] backbone to extract task-specific features in each task branch. To share information between two branches and benefit from the interactions of these two task-specific features concerning the nature of the Transformer's globality, $T^2$Net uses a unique Transformer design to learn a generalized representation. Since the reconstruction branch has more potential in artifact removal capacity than the super-resolution branch, the task Transformer module guides the super-resolution branch into high-quality representation from the reconstruction branch. The Transformer module inherits the query ($Q$: from super-resolution branch), key ($K$: from reconstruction branch), and value ($V$: from reconstruction branch) from each scale's two branches' output. It forms three main concepts: relevance embedding, Transfer attention, and soft attention, which differ from the original Transformer blocks. Relevance embedding tries to enclose the correlated features from the reconstruction branch to the super-resolution branch. Transfer attention aims to transmit the anatomical and global features between two branches, and last but not least, soft attention amalgamates features from the previous two steps. Ultimately, this module lets the whole net-

work transfer and synthesize the representative and anatomical features to produce a high-quality, artifacts-free representation from highly undersampled measurements. The experimental results on two datasets expressed the high potential of this approach rather than conventional algorithms.
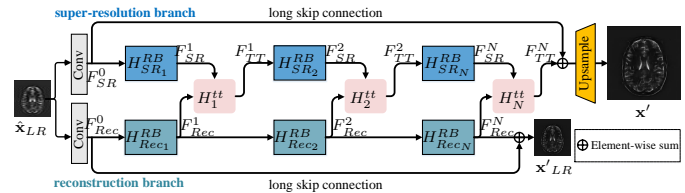


(a) An overview of $T^2$Net [272], a multi-task learning framework that consists of a super-resolution branch and reconstruction branch. The reconstruction branch embraces the stronger capability of artifact removal therefore, the task Transformer module is fed with the reconstruction branch.
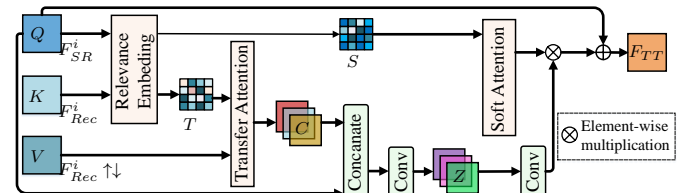


(b) inner design of proposed $H^{tt}$ task Transformer module. Analogous to Figure 2, the design of $T^2$ module follows the naive design with some modifications, and in contrast to seminal design, all $Q$, $K$, and $V$ entities do not originate from the same representation—$Q$ comes from the super-resolution branch and the rest from the reconstruction branch.

Figure 29: An overview of $T^2$Net [272]. (a) Multi-Task $T^2$Net pipeline and (b) Task Transformer Module—$T^2$ Module

### 5.4. Super-Resolution Reconstruction

Improving the resolution of images leads to the more detailed delineation of objects. Increasing the medical image resolution plays a crucial role in computer-aided diagnosis due to its rich anatomical and textural representation. Based on the aforementioned fact and the MRI's pipeline physics during the image acquisition process for having high-resolution images, a patient needs to lie a long time in the MRI tube. Hereupon

Table 9: Medical Image Reconstruction. *LDE*, *SVR*, *USR*, and *SRR* stand for Low Dose Enhancement, Sparse-View Reconstruction, Undersampled Reconstruction, and Super Resolution Reconstruction, respectively. † indicates that this network uses a pre-trained perceptual loss (loss network).

| Method | Task(s) | Modality | Type | Pre-trained Module: Type | Dataset(s) | Metrics | Year |
|---|---|---|---|---|---|---|---|
| **Pure** | | | | | | | |
| TED-net [226] | LDE | CT | 2D | ✗ | NIH-AAPM-Mayo Clinical LDCT [252] | SSIM RMSE | 2021 |
| Eformer [227] | LDE | CT | 2D | ✗† | NIH-AAPM-Mayo Clinical LDCT [252] | PSNR, SSIM RMSE | 2021 |
| CTformer [230] | LDE | CT | 2D | ✗ | NIH-AAPM-Mayo Clinical LDCT [252] | SSIM RMSE | 2023 |
| FIT [233] | SVR | CT | 2D | ✗ | LoDoPaB [261] | PSNR | 2021 |
| ViT-Rec [236] | USR | MRI | 2D | Supervised | fastMRI [262] | SSIM | 2021 |
| **Encoder** | | | | | | | |
| TransCT [225] | LDE | CT | 2D | ✗ | [1] NIH-AAPM-Mayo Clinical LDCT [252] [2] Private clinical pig head CBCT | RMSE SSIM VIF | 2021 |
| STFNet [229] | LDE | PET MRI | 2D | ✗ | Private Dataset | RMSE, PSNR SSIM, PCC | 2022 |
| SIST [231] | LDE | CT | 2D | ✗ | [1] LDCT Dataset [263] [2] Private dataset | PSNR, SSIM RMSE | 2022 |
| DuDoTrans [232] | SVR | CT | 2D | ✗ | NIH-AAPM-Mayo Clinical LDCT [252] | PSNR, SSIM RMSE | 2021 |
| CTTR [234] | SVR | CT | 2D | ✗ | LIDC-IDRI [264] | RMSE, PSNR SSIM | 2022 |
| **Skip Connection** | | | | | | | |
| Cohf-T [239] | SRR | MRI | 2D | ✗ | [1] BraTS2018 [202] [2] IXI [265] | PSNR SSIM | 2022 |
| **Other Architectures** | | | | | | | |
| 3D T-GAN [228] | LDE | PET | 3D | ✗ | Private Dataset | PSNR, SSIM NMSE | 2021 |
| ARMLUT [235] | SVR | CT | 3D | ViT: Supervised † | [1] SPARE Challenge Dataset [266] [2] Walnut dataset [267] | PSNR, SSIM | 2022 |
| $T^2$Net [237] | USR SRR | MRI | 2D | ✗ | [1] IXI [265] [2] Private Dataset | PSNR, SSIM NMSE | 2021 |
| DisCNN-ViT [238] | SRR | MRI | 3D | ViT: Self-Supervised | [1] fastMRI [262] [2] IXI [265] | PSNR, SSIM NMSE | 2021 |
| SLATER [240] | USR | MRI | 2D | ViT: Self-Supervised | [1] fastMRI [262] [2] IXI [265] | PSNR, SSIM | 2022 |
| DSFormer [241] | USR | MRI | 2D | ViT: Self-Supervised | IXI [265] | PSNR, SSIM | 2023 |

lower signal-to-noise ratio and more minor spatial coverage drawbacks are inevitable [269]. Therefore in this section, we investigate Transformer-utilized algorithms that try to alleviate this problem. Of note, due to the analogy between MRI and super-resolution reconstruction, some studies investigate these two tasks in conjunction with each other.

Mahapatra et al. [238] proposed the GAN-based model with structural and textural information preservation done by multiple loss function terms. Their pipeline included two pre-trained modules named feature disentanglement module, a conventional autoencoder, and a Transformer-based feature encoder, UNETR [159]. UNETR captures the global and local context of the original low-resolution image and induces the high-resolution image to preserve these contexts too. These two modules fine-tune on a different medical dataset, and afterward, the low-resolution input plus the intermediate generator produced image feed to these modules. The disentanglement

network contains two autoencoders to learn two counterparts, latent space, structural and textural space, with fed medical images. In an end-to-end setting, these two pre-trained assistive modules help to generate more realistic and structural and textural preserving high-resolution images by imposing module-related loss terms such as adversarial loss to constrain for producing realistic images and cosine similarity loss for each mentioned module. Results on the IXI dataset proved that Mahapatra et al.'s [238] network outperformed a couple of the CNN-based attention mechanism networks and $T^2$Net [237]. Maintaining structural information during the acquiring high-resolution images plays a crucial role. Hence, the structure information is embedded in an image's high-frequency counterpart, like in an image's gradients. In addition, due to the less time-consuming nature of obtaining MR T1WI (T1 Weighted Image), it is wise to use it as an inter-modality context prior to producing a high-resolution image. Accordingly, Fang et

Table 10: A brief description of the reviewed Transformer cooperated in the medical image reconstruction field.

| Method | Contributions | Highlights |
|---|---|---|
| **Low Dose Enhancement** | | |
| TransCT [225] | • The proposed prototype was the first successful implementation of a Transformer complement to CNN in the Low Dose CT reconstruction domain by exploring its revenue within high-frequency and low-frequency counterparts. | • The Transformer effectively could learn the embedded texture representation from the noisy counterpart.<br>• This paradigm is not convolution-free and uses Transformer as a complement. |
| TED-net [226] | • Convolution-free U-Net based Transformer model. | • Introduced Dialted Token-to-Token-based token serialization for an improved receptive field in Transformers.<br>• Using Cyclic Shift block for feature refinement in tokenization. |
| Eformer [227] | • Incorporate the learnable Sobel filters into the network for preserving edge reconstruction and improving the overall performance of the network.<br>• Conduct diverse experiments to validate that the residual learning paradigm is much more effective than other learning techniques, such as deterministic learning approaches. | • Successfully imposed the Sobel-Feldman generated low-level edge features with intermediate network layers for better performance.<br>• To guarantee the convergence of the network, they used the Multi-scale Perceptual (MSP) loss alongside Mean Squared error (MSE) to hinder the generation of disfavored artifacts. |
| 3D T-GAN [228] | • It is a 3D-based method rather than conventional 2D methods.<br>• First LDPET enhancement study that leveraged from Transformer to model long-range contextual information. | • To produce more reliable images with a generator they used an adversarial loss to make the data distribution the same as real data. |
| STFNet [229] | • Proposed a dual input U-Net-based denoising structure for low-count PET images with excessive MRI modality contribution.<br>• Used a Transformer block as a hybrid add-on for feature fusion to make a pixel-to-pixel translation of PET and MRI modalities. | • In comparison with the U-Net and residual U-Net structures due to the different training strategy which is roughly named Siamese structure has a low computational burden and simplified network.<br>• This network successfully handled the disparity and nonuniformity of shape and modality of PET and MRI.<br>• The visual results of denoised images testify that the proposed method could recover the detail of texture more clearly than other networks. |
| CTformer [230] | • Convolution-free, computational efficient design.<br>• Introduce a new inference mechanism to address the boundary artifacts.<br>• Proposed interpretability method to follow each path resultant attention map through the model to understand how the model is denoising. | • Alleviate receptive filed deficiency with the token rearrangement. |
| SIST [231] | • Proposed inner-structure loss to mimic the physics of the functionality of sinogram processing by CT devices to restrain the noise.<br>• Extracting the long-range dependencies between distinct sinogram angles of views via Transformer. | • Utilizing the image reconstruction module to alleviate the artifacts that could happen in sinogram domain denoising by transferring the sinogram noise into the image domain.<br>• Image domain loss back-propagates into the sinogram domain for complementary optimization. |
| **Sparse-View Reconstruction** | | |
| DuDoTrans [232] | • To cope with the global nature of the sinogram sampling process introduced, the SRT module, a hybrid Transformer-CNN to capture long-range dependencies.<br>• Utilizing a dual domain model to simultaneously enrich raw sinograms and reconstruct CT images with both enhanced and raw sinograms.<br>• To compensate for the drift error between raw and enhanced sinogram representation employs DuDo Consistency Layer. | • Utilizing a residual learning paradigm for image-domain reconstruction.<br>• Fewer parameters in comparison with other structures, e.g., DuDoNet [268] with better performance. |
| FIT [233] | • Introduced the Fourier Domain Encoding to encode the image to a lower resolution representation for feeding to the encoder-decoder Transformer for reconstructing the sparse-view CT measurements. | • Introduced the Fourier coefficient loss as a multiplicative combination of amplitude loss and phase loss in the complex domain. |
| CTTR [234] | • Introduced the encoder-decoder Transformer pipeline to utilize dual-domain information, raw sinogram, and primary reconstruction of CT via FBP [256] algorithm for sparse-view CT measurement reconstruction. | • In contrast to DuDoTrans [232], CTTR directly utilizes raw sinograms to enhance reconstruction performance. |
| ARMLUT [235] | • Proposed a paradigm for CT image reconstruction in a non-trainable manner.<br>• Extending the most DIP research on 2D to 3D medical imaging scenario. | • Optimising the large-scale 3D Transformer with only one reference data in an unsupervised manner.<br>• Stabilising the iterative optimization of multi-loss untrained Transformer via re-weighting technique. |
| **Undersampled Reconstruction** | | |
| ViT-Rec [236] | • This study investigated the advantage of pure Transformer framework, ViT, in fastMRI reconstruction problems in comparison with the baseline U-Net.<br>• ViT benefits from less inference time and memory consumption compared to the U-Net. | • Utilizing pre-training weights, e.g., ImageNet, extensively improves the performance of ViT in the low-data regime for fastMRI reconstruction, a widespread concept in the medical domain.<br>• ViTs that accompany pre-training weights demonstrate more robust performance toward anatomy shifts. |
| T²Net [237] | • Introduce the first Transformer utilized multi-task learning network in the literature.<br>• Designed the task Transformer for maintaining and feature transforming between branches in the network.<br>• Outperformed the sequentially designed networks for simultaneous MRI reconstruction and super-resolution with T²Net. | • Used the same backbone for feature extraction in branches, however, the purpose of the branches is diverse. |
| **Super Resolution Reconstruction** | | |
| DisCNN-ViT [238] | • Using a Transformer-based network to capture global contextual cues and amalgamate them with CNN's local information results in the superior quality of high-resolution images in super-resolution literature.<br>• Creating realistic images is just not a burden on an adversarial loss function, in addition, multiple loss functions incorporate extra constraints that preserve anatomical and textural information in the begotten image. | • Multi prerequisite steps are required to train the experiments; however, the super-resolution step is a straightforward end-to-end network.<br>• Need fine-tuning steps for two disentanglement and UNETR networks.<br>• The computational burden of UNETR is high and could use new efficient transformer-designed networks. |
| Cohf-T [239] | • Leverage the high-resolution T1WI due to its rich structural information for super-resolving T2-weighted MR images.<br>• Introduced the high-frequency structure prior and intra-modality and inter-modality attention paradigms within the Cohf-T framework. | • Assess prior knowledge into super-resolution paradigm successfully.<br>• Competitive number of FLOPS in reaching SOTA PSNR results in comparison with other attention-based networks.<br>• End-to-end pipeline for training the network. |
| **Self-supervised Reconstruction** | | |
| SLATER [240] | • A Self-supervised MRI reconstruction framework that uses unconditional adversarial network and cross-attention transformers for high-quality zero-shot reconstruction.<br>• The first work that proposes an adversarial vision transformer framework designed for the task of MRI reconstruction. | • SLATER outperforms CNN and self-attention GAN models, showing better invertibility and improved reconstruction performance with cross-attention transformer blocks and unsupervised pretraining.<br>• Prior adaptation is computationally burdensome but improves across-domain reconstructions. |
| DSFormer [241] | • Proposes a deep learning model for fast MRI reconstruction that uses a deep conditional cascade transformer and self-supervised learning to exploit contrast redundancy and reduce data requirements. | • Both k-space and image-domain self-supervision independently produce strong reconstruction quality, showing the benefits of self-supervision in both domains. |

al. [239] devised a network (see Figure 30) to leverage these two concerns in their super-resolution pipeline: Cross-Modality High-Frequency Transformer (**Cohf-T**). This network is divided into two streams, the first stream is applied on low-resolution T2WI, and the second one manipulates T2WI's gradient and high-resolution T1WI. The Cohf-T module interacts between two streams to embed the prior knowledge in the super-resolution stream's features. The Cohf-T module consists of three different attention modules: short-distance and long-distance window attention and inter-modality attention. The first two attention modules help to model intra-modality dependency. To be precise, the short-distance window helps recover the local discontinuity in boundaries with the help of surrounding structure information, and the long-distance window can capture the textural and structural patterns for enhanced re-

sults. Due to the discrepancy in intensity levels between T1WI and T2WI, it is vital to make an alignment between these two domains, and Fang et al. [239] presented a Feature Alignment (FA) module to reduce the cross-modality representation gap. They compared their results with T²Net [237] and MTrans [273], which outperformed both approaches by ∼ 1% in terms of PSNR.

## 5.5. Self-supervised Reconstruction

Recently, self-supervised reconstruction methods requiring only undersampled k-space data have been proposed for single-contrast MRI reconstruction, enabling the autonomous learning of high-quality image reconstructions without the need for fully sampled data or manual annotations. **SLATER** [240] is introduced as a self-supervised MRI reconstruction method
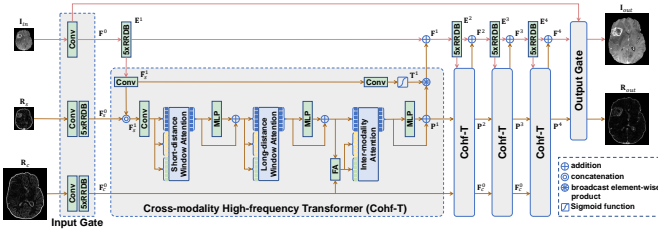
Figure 30: The pipeline of Cohf-T [239] consists of three main branches with the corresponding input modalities as follows: $\mathbf{I}_{in}$, $\mathbf{R}_s$, and $\mathbf{R}_c$ denote the low-resolution T2WI, the gradient of low-resolution T2WI and high-resolution T1WI, respectively. A fully-convolutional branch for density-domain super-resolution, a Transformer-based branch for restoring high-frequency signals in the gradient domain, and a guidance branch for extracting priors from the T1 modality. *Conv*, *RRDB* and *MLP* represent a $3 \times 3$ convolution operation and residual-in-residual dense block and multi-layer perceptron, respectively.

that combines an unconditional adversarial network with cross-attention transformers. It learns a high-quality MRI prior during pre-training and performs zero-shot reconstruction by optimizing the prior to match undersampled data. SLATER improves contextual feature capture and outperforms existing methods in brain MRI. It has potential for other anatomies and imaging modalities, offering reduced supervision requirements and subject-specific adaptation for accelerated MRI. Likewise, **DSFormer** [241] is presented for accelerated multi-contrast MRI reconstruction. It uses a deep conditional cascade transformer (DCCT) with Swin transformer blocks. DSFormer is trained using self-supervised learning in both k-space and image domains. It exploits redundancy between contrasts and models long-range interactions. It also reduces the need for paired and fully-sampled data, making it more cost-effective.

### 5.6. Discussion and Conclusion

In this section, we outline different Transformer-based approaches for medical image reconstruction and present a detailed taxonomy of reconstruction approaches. We reviewed 17 studies that profit from the Transformer design to compensate for the deficiency of CNN's limited receptive field, with two of them incorporating self-supervised strategies [240, 241] for training transformer modules. We investigate each study in depth and represent Table 9 for detailed information about the dataset, utilized metrics, modality, and objective tasks. In Table 10, we provide the main contribution of each study and the prominent highlight of each method, and in Table 8, we present a comparison of various methods for the low-dose enhancement task using the NIH-AAPM-Mayo [252] benchmark medical reconstruction dataset, which shows that Eformer [227] is superior as compared to TransCT [225] and CTformer [230], considering its higher SSIM and lower RMSE scores.

Most of the studies in this domain use the original Transformer as a plug-and-play module in their design and only a limited number of studies utilize hierarchical and efficient Transformers. However, the criteria for using multi-scale and hierarchical architectures are generally important for dense prediction tasks, e.g. image reconstruction, and

should be considered further. Also, another direction to follow for future research could be to investigate the influence of using pre-training weights on Transformers due to the need for a large amount of data for better convergence results in Transformers, which contradicts the nature of the medical domain, due to the scarceness of annotated medical data.

In addition, we noticed that most of the studies focus on MRI and CT image reconstruction tasks. So there is a need for evaluating the applicability of these methods on other modalities, too.

## 6. Medical Image Synthesis

In this section, we will overview several instances of Transformers in the medical image synthesis task. The scarcity of medical data and the high cost of acquisition processes make this task very valuable in the medical field. Some studies aim to synthesize missing slices from MRI and CT sequences. In addition, some methods target capturing the structural information in diverse modalities, e.g., CT to MRI image-to-image translation and vice versa. Figure 31 shows our taxonomy for the image-synthesized methods.
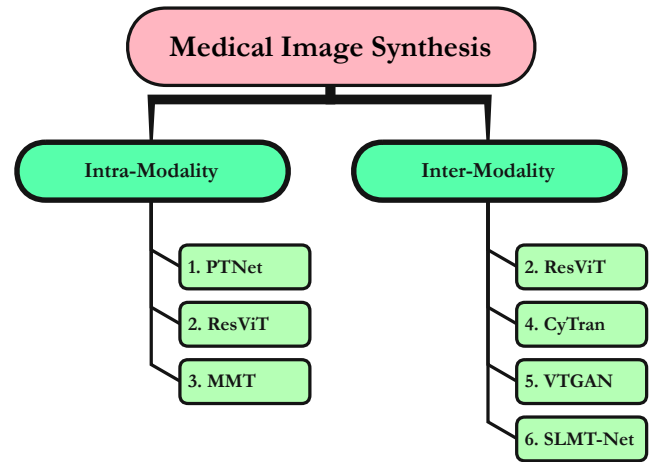


Figure 31: An overview of ViTs in medical image synthesis. Methods are categorized by target and source modality. The prefix numbers in the paper's name in ascending order denote the reference for each study as follows: 1. [274], 2. [275], 3. [276], 4. [277], 5. [278], 6. [279].

### 6.1. Intra-Modality

The main objective of the intra-modality methods is to synthesize high-quality images using low-quality samples from the same modality. In this respect, several Transformer-based approaches are presented to formulate the synthesis task as a sequence-to-sequence matching problem to generate fine-grained features. In this section, we will briefly present some recent samples [274].

Brain development monitoring is a de facto standard in predicting later risks; hence it is critical to screen brain biomarkers via available imaging tools from early life stages. Due to

this concern and the nature of MRI and subjective infants' restlessness, it is not relatively straightforward to take all the MR modalities during the MRI acquisition. Zhang et al. [274] proposed a **Pyramid Transformer Net (PTNet)** as a tool to reconstruct realistic T1WI images from T2WI. This pipeline is an end-to-end Transformer-based U-Net-like and multi-resolution structure network utilizing an efficient Transformer, Performer [280], in its encoder (PE) and decoder (PD). Analogously to the original U-Net [28], they used skip connection paths for preserving fine-grained features and accurate localization features for reconstruction. Moreover, the paradigm's two-level pyramidal design helps the network capture local and global information in a multi-resolution fashion. They achieved the SOTA results on the dHCP [281] dataset compared with the flagship GAN-based image generation method pix2pix (HD) [282, 283].
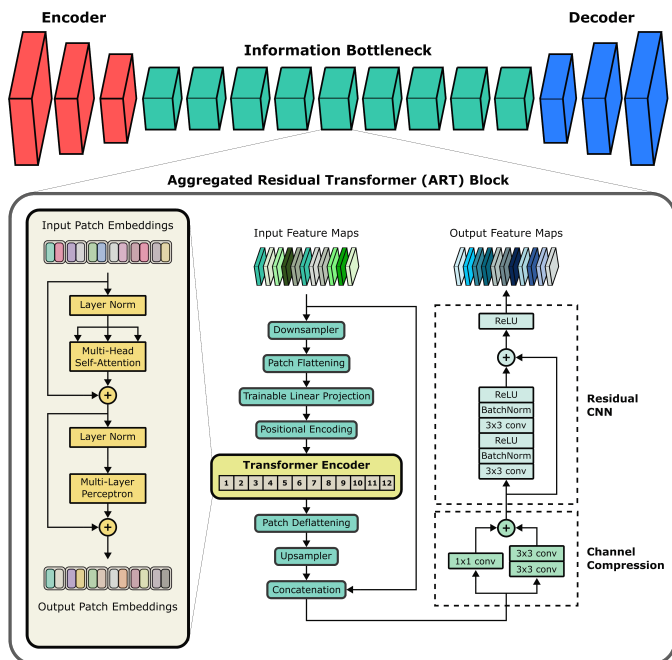


Figure 32: The ResViT [275] framework for multi-modal medical image synthesis. The bottleneck of this encode-decoder comprises successively residual Transformers and residual convolutions layers for synergistically capturing the fine-grained global and local context.

Dalmaz et al. [275] introduced a conditional generative adversarial network based on the cooperation of Transformers and CNN operators, namely **ResViT**. This paradigm addresses the issue of needing to rebuild separate synthesis models for varying source-target modality settings and represents a unified framework as a single model for elevating its practicality. The ResViT (Figure 32) pervasively refers to the generator of its pipeline, whereby it leverages a hybrid pipeline of residual convolutional operations and Transformer blocks that enable effective aggregation of local and long-range dependencies. The discriminator is based on a conditional PatchGAN framework [282]. Utilizing standalone Transformer architectures (e.g., PTNet [274]) in pixel-to-pixel tasks is quite challenging due to the quadratic complexity, which limits its usage to fixed-size patches that hamper its effectiveness. From Figure 32, it is evident that residual Transformer blocks stacked successively,

known as aggregated residual Transformer (ART) blocks, in the bottleneck of the encoder-decoder design of the generator to extract the hidden contextual information of input features. The primary motivation of ART blocks is to learn an integrated representation that combines contextual, local, and hybrid local-contextual features underhood from the input flow. Channel Compression (CC) module recalibrates the concatenated features from the previous ART block and Transformer module to select the most discriminant representations. Due to the cascade of Transformers in design, to decrease the model complexity and computational burden, ResViT utilizes weight sharing strategy among projection tensors for Query, Key, value, and attention heads besides weight matrices for multi-layer perceptron operation. The superiority of this method has been proved over several MRI datasets in multi-contrast MRI synthesis and MRI to CT experiments with high PSNR and SSIM metrics over the conventional SOTA methods, e.g., pGAN [284], SAGAN [285], pix2pix [282] and PTNet [274].

Likewise, Liu et al. [276] addressed the issue of missing contrasts in MRI imaging and proposed a multi-contrast multi-scale Transformer (**MMT**) framework to handle the unavailability of this information by synthesizing the existing contrasts as a means to substitute the missing data. To achieve efficient contrast synthetization, the task is considered as a seq-to-seq problem, in which the model learns to generate missing contrasts by leveraging the existing contrast in the following manner: A Swin multi-contrast Transformer encoder is implemented that creates hierarchical representation from the input MRI image. Then, a Swin Transformer-based architecture decodes the provided representation at multiple scales to perform medical image synthesis. Both the encoder and decoder are composed of two sequential swin blocks that capture contrast dependencies effectively. Conducted experiments on the IXI [265] and BraTS [206] datasets demonstrated MMT's advantage compared to previous methods.

Table 11: Comparison results on BraTS [206] dataset for medical image synthesis, focusing on one-to-one and many-to-one tasks with different input to output combinations of T1, T2, and FLAIR modalities. The symbol (†) indicates that the values are taken from [279]. The GFLOPs of all methods are calculated with an input image size of 256×256.

| Methods | Design | Task | SSIM ↑ | PSNR ↑ | Params (M) | FLOPs (G) |
|---|---|---|---|---|---|---|
| PTNet † [279] | Hybrid | T1, T2 → FLAIR | 0.851 | 23.78 | 27.69 | 233.1 |
|  |  | T1, FLAIR → T2 | 0.905 | 25.09 |  |  |
|  |  | T2, FLAIR → T1 | 0.920 | 22.19 |  |  |
|  |  | T2 → FLAIR | 0.851 | 23.01 |  |  |
|  |  | FLAIR → T2 | 0.894 | 24.78 |  |  |
| MMT [276] | Pure | T1, T2 → FLAIR | 0.909 | 26.20 | 138.2 | 221.5 |
|  |  | T1, FLAIR → T2 | 0.934 | 26.74 |  |  |
|  |  | T2, FLAIR → T1 | 0.922 | 25.43 |  |  |
|  |  | T2 → FLAIR | 0.891 | 25.00 |  |  |
|  |  | FLAIR → T2 | 0.902 | 24.84 |  |  |
| ResViT [275] | Bottleneck | T1, T2 → FLAIR | 0.886 | 25.84 | 123.4 | 486.1 |
|  |  | T1, FLAIR → T2 | 0.938 | 26.90 |  |  |
|  |  | T2, FLAIR → T1 | 0.924 | 26.20 |  |  |
|  |  | T2 → FLAIR | 0.870 | 24.97 |  |  |
|  |  | FLAIR → T2 | 0.908 | 25.78 |  |  |

## 6.2. Inter-Modality

Unlike the intra-modality strategies, the inter-modality methods are designed to learn the mapping function between two different modalities. This approach allows the network to convert the samples from the base modality into a new modality

and leverage the generated samples in the training process for the sake of performance gain. In this section, we will elaborate on two Transformer-based [277, 278] strategies.

Several medical conditions may prevent patients from receiving intravenous contrast agents while getting CT screening. However, the contrast agent is crucial in assisting medical professionals in identifying some specific lesions. Therefore, **CyTran** [277] is proposed as an unsupervised generative adversarial convolutional Transformer for translating between contrast and non-contrast CT scans and image alignment of contrast-enhanced CT scans to non-enhanced. Its unsupervised part is also derived from its cyclic loss. CyTran is composed of three main modules: **I)** A downsample CNN-based module designed for handling high-resolution images, **II)** A convolutional Transformer module tailored for incorporating both local and global features, and **III)** An upsampling module developed to revert the transformation of the downsampling block through transpose-convolution. Additionally, the authors introduce a new dataset, Coltea-Lung-CT-100W, comprised of 100 3D anonymized triphasic lung CT scans of female patients.

Furthermore, Kamran et al. [278] trained a ViT-based generative adversarial network (**VTGAN**) in a semi-supervised fashion on the Fundus Fluorescein Angiography (FFA) dataset provided by [288] via incorporating multiple modules, including residual blocks as generators for coarse and fine image generation, and two ViT architectures consisting of identical transformer encoder blocks for concurrent retinal abnormality classification and FA image synthesis.

To embed self-supervised training in medical image synthesis, **SLMT-Net** [279] is presented for generating missing modalities in magnetic resonance (MR) images. It tackles the limited paired data challenge by combining paired and unpaired data. The method involves edge-aware pre-training using an Edge-preserving Masked Auto-Encoder (Edge-MAE) model that predicts missing information and estimates the edge map. A patch-wise loss enhances the performance. In the multi-scale fine-tuning stage, a DSF module is integrated to synthesize missing-modality images using multi-scale features. The pre-trained encoder extracts high-level features to ensure consistency between synthesized and ground-truth images during training.

### 6.3. Discussion and Conclusion

This section covers the adoption of ViT architectures in medical image synthesis applications. We explored the proposed methods based on two synthesis approaches: **(I)** inter-modality, in which the target modality is synthesized in a way that encapsulates crucial diagnostic features from different source images; and **(II)** intra-modality, with the objective of yielding target images with better quality by integrating information from lower resolution source images. To demonstrate their effectiveness, these approaches usually rely on SSIM, PSNR, and LPIPS as the evaluation metrics, since they are designed to measure the similarity between images. We also reviewed a ViT-based synthesis model [278] that operates in a decoder fashion for the

task of fundus-to-angiogram translation with different evaluation measurements, including Fréchet Inception Distance (FID) and Kernel Inception Distance (KID). A combination of self-supervised learning with transformer-based medical synthesis is further reviewed to emphasize the cost-saving benefits of reduced labeling [279]. We have additionally provided the architectural type, modality, input size, training setting, datasets, metrics, and year for every medical registration technique analyzed in Table 12. Furthermore, Table 13 lists the contributions and highlights of the proposed works. To facilitate comparison, table Table 11 includes performance evaluation of multiple methods on the BraTS dataset [206]. In terms of SSIM and PSNR scores, MMT [276] typically achieves superior performance, while PTNet [279] falls slightly behind when compared to the other two methods. Moreover, we have provided information about the model size and computational cost for three distinct approaches: PTNet, MMT, and ResViT. PTNet, with 27.69 million parameters and 233.1 GFLOPs, strikes a balance between model size and computational demand. Meanwhile, MMT offers a larger model with 138.2 million parameters but lower computational requirements at 221.5 GFLOPs, emphasizing efficiency. On the other hand, ResViT provides a mid-sized model with 123.4 million parameters but a notably higher computational load of 486.1 GFLOPs. The choice among these methods should be guided by specific task requirements, weighing the trade-off between model size and computational cost to achieve optimal performance. However, it is important to note that all these models require a significant computational cost.

With the scarcity of works with ViT implementations and the recent advancement in the medical synthesis field with Transformer models, we believe that these systems require more research effort to be put into them. For example, Transformers have much room for improvement to generate more realism and high-quality synthesized medical images. One way to achieve this is by incorporating more detailed anatomy and physiology features using more efficient and effective attention mechanisms. Additionally, while much of the current research in this area has focused on 2D medical images and CT and MRI modalities, there is potential to apply these techniques to other types of medical images, including 3D and microscopy images.

## 7. Medical Image Detection

Object detection remains one of the challenging problems in computer vision, especially detection in the medical image domain has its own challenges. Current SOTA architectures that work on 2D natural images use Vision Transformers. The Vision Transformers used in the detection task can be classified into two Transformer backbones and detection Transformers. In addition, the Transformer module can be used in a hybrid manner. Detection Transformers generally represent an end-to-end detection pipeline with an encoder-decoder structure, while the Transformer backbone solely utilizes the Transformer en-

Table 12: An overview of the reviewed Transformer-based medical image synthesizing approaches.

| Method | Concept(s) | Modality | Type | Pre-trained Module: Type | Dataset(s) | Metrics | Year |
|--------|-----------|----------|------|--------------------------|-----------|---------|------|
| **Pure** | | | | | | | |
| PTNet [274] | Intra-Modality | MRI | 2D | ✗ | dHCP dataset [281] | SSIM PSNR | 2021 |
| MMT [276] | Intra-Modality Inter-Modality | MRI | 2D | ✗ | [1] IXI [265] [2] BraTS [206] | SSIM PSNR LPIPS | 2022 |
| **Bottleneck** | | | | | | | |
| ResViT [275] | Intra-Modality Inter-Modality | CT MRI | 2D | ViT: Supervised | [1] IXI [265] [2] BraTS [200, 201, 202] [3] Multi-modal pelvic MRI-CT [286] | PSNR SSIM | 2022 |
| CyTran [277] | Inter-Modality | CT | 2D 3D | ✗ | Coltea-Lung-CT-100W [277] | MAE RMSE SSIM | 2022 |
| **Encoder** | | | | | | | |
| SLMT-Net [279] | Inter-Modality | MRI | 2D | ViT: Self-supervised | BraTS [200, 201, 202] [2] ISLES [287] | SSIM PSNR NMSE | 2022 |
| **Decoder** | | | | | | | |
| VTGAN [278] | Inter-Modality | Angiography | 2D | ✗ | Fundus Fluorescein Angiography [288] | Fréchet inception distance Kernel Inception distance | 2021 |

Table 13: A brief description of the reviewed Transformer-based medical image synthesizing models.

| Method | Contributions | Highlights |
|--------|--------------|-----------|
| **Pure** | | |
| PTNet [274] | • Introduced the pure Transformer-based network with linear computational complexity for image-synthesizing context | • Practical inference time around 30 image/s |
| MMT [276] | • Proposed a pure Transformer-based architecture that incorporates Swin Transformer blocks to perform missing data imputation by leveraging the existing MRI contrasts. • Conducted experiments on the IXI and BraTS datasets to perform qualitative and quantitative analysis and confirm their model's efficiency. | • Since the attention mechanism can be utilized to pinpoint influential features in the model's reasoning and decision-making, the attention scores of the Transformer decoder in MMT make it interpretable by capturing information in different contrasts that play an important role in generating the output sequence. • The framework can be applied to a variety of medical analysis tasks, including image segmentation and cross-modality synthesis. |
| **Bottleneck** | | |
| ResViT [275] | • First conditional adversarial model for medical image-to-image translation with hybrid CNN-Transformer generator • Introduced a new module, ART block, for simultaneously capturing localization and contextual information | • Utilized weight sharing strategy among Transformers to hinder the computational overhead and lessen the model complexity • An end-to-end design for the synthesized model that generalizes through multiple settings of source-target modalities, e.g., one-to-one and many-to-one tasks |
| CyTran [277] | • Proposing a generative adversarial convolutional Transformer for two tasks of image translation and image registration. • Introducing a new dataset, named Coltea-Lung-CT-100W, comprised of 100 3D anonymized triphasic lung CT scans of female patients. | • The presented method can handle high-resolution images due to its hybrid structure • Utilized style transfer techniques to improve alignment between contrast and non-contrast CT scans |
| **Encoder** | | |
| SLMT-Net [279] | • Proposed a self-supervised method for cross-modality MR image synthesis to generate missing modalities based on available modalities. • Addresses the challenge of limited paired data by combining both paired and unpaired data. | • Despite using only 70% paired data, MT-Net's performance remains on par with comparable models. • Yields Enhanced PSNR scores when utilized for MR image synthesis. • The combination of edge-aware pre-training and multi-scale fine-tuning stages contributes to the success of SLMT-Net. |
| **Decoder** | | |
| VTGAN [278] | • Proposed a synthesis model for the task of fundus-to-angiogram that incorporates ViT architecture in the decoder section of the system to concurrently classify retinal abnormalities and synthesize FA images. • Prepared experimental data based on quantitative and qualitative metrics regarding the model's generalization ability under the influence of spatial and radial transformations. | • Has the potential to be adopted as a tool for tracking disease progression. • The system is designed to operate on non-invasive and low-cost fundus data to generate FA images. |

coder for feature refinement. In order to increase detection performance, object detectors combine variants of vision Transformers with classical convolutional neural networks (CNNs).

Quite recently, Carion et al. [187] introduced the concept of DETR, which forms a foundation for Detection Transformers. DETR uses a ResNet backbone to create a lower-resolution representation of the input images. Even though this approach achieves very good 2D detection results, comparable to the R-CNN backbone, high computational complexity is a downside of this method. The deformable DETR [16] approach has improved DETR's detection performance overcoming the problem of high computational complexity. Many recent approaches have tried to improve DETR's detection concept over time. Efficient DETR [297] eliminated DETR's requirement for iterative refinement. Conditional DETR [298] introduced the concept of a conditional cross-attention module. DN-DETR [299] introduced a denoising strategy, and DINO [300] im-

proved many aspects, such as denoising training, etc. Recently, some studies performed experiments on 2D medical data such as [301], [293], etc. However, only very few attempts tried to adapt it to 3D. Spine-Transformer was proposed by Tao et al. [295] for sphere-based vertebrae detection. Another approach in 3D detection was proposed by Ma et al. [289], which introduced a novel Transformer that combines convolutional layers and Transformer encoders for automatically detecting coronary artery stenosis in Coronary CT angiography (CCTA). An approach to better extract complex tooth decay features was proposed by [290]. For end-to-end polyp detection, Shen et al. [293] proposed an approach that was based on the DETR model. Kong et al. have proposed the approach CT-CAD [294], context-aware Transformers for end-to-end chest abnormality detection. As illustrated in Figure 33, we have broadly classified these methods based on the role ViT plays in the detection's commonly accepted pipeline. Table 14 indicates details
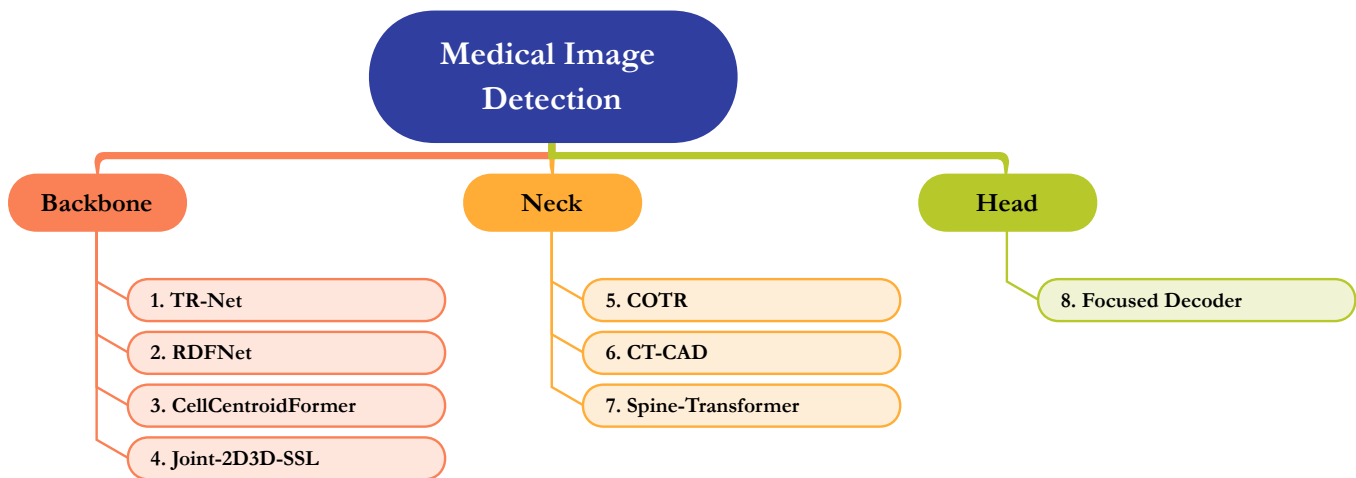
Figure 33: An overview of Transformers in medical image detection. Methods are classified into the backbone, neck, and head according to the positions of the Transformers in their architecture. The prefix numbers in the paper's name in ascending order denote the reference for each study as follows: 1. [289], 2. [290], 3. [291], 4. [292], 5. [293], 6. [294], 7. [295], 8. [296].

on modalities, organs, datasets, metrics, etc. The highlights of different approaches are summarized in Table 15. Some of the aforementioned detection papers in the medical image domain are reviewed in detail in this section.

### 7.1. Backbone

This section explains Transformer networks using only the Transformer encoder layers for object detection. The work proposed by Ma et al. [289] uses a Transformer network **(TR-Net)** for identifying stenosis. A leading threat to the lives of cardiovascular disease patients globally is Coronary Artery Disease (CAD). Hence, the automatic detection of CAD is quite significant and is considered a challenging task in clinical medicine. The complexity of coronary artery plaques, which result in CAD, makes the detection of coronary artery stenosis in Coronary CT angiography (CCTA) challenging.

The architecture introduces a Transformer and combines the feature extraction capability of convolutional layers and Transformer encoders. TR-Net can easily analyze the semantic information of the sequences and can generate the relationship between image information in each position of a multiplayer reformatted (MPR) image. This model can effectively detect stenosis based on both local and global features. The CNN easily extracts the local semantic information from images, and the Transformer captures global semantic details more easily. A 3D-CNN is employed to capture the local semantic features from each position in an MPR image. After this step, the Transformer encoders are mainly used to analyze feature sequences. The main advantage here is that this helps in mining the dependency of local stenosis on each position of the coronary artery. The architecture of the TR-Net is given in Figure 34. One part of the figure indicates the 3D-CNN. This module extracts the local features. The other part indicates the Transformer encoder structure. This module associates the local feature maps of each position and helps in analyzing the dependency between different positions, which in turn is helpful for classifying the significant stenosis at each position. The CNN part mainly has two

main advantages: it prevents the overfitting of semantic information and improves the model's efficiency. The input to the network architecture is the coronary artery MPR image.
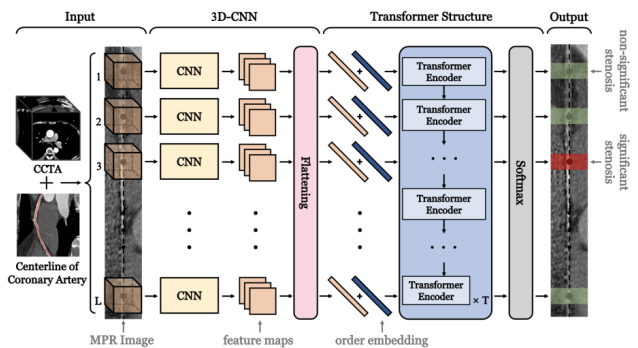


Figure 34: Proposed architecture of TR-Net model [289].

The 3D-CNN module has four sequentially connected substructures, which consist of a convolutional kernel of size $3 \times 3 \times 3$, a non-linear ReLU layer and a $2 \times 2 \times 2$ max-pooling layer. The number of filters is 16 in the first part, and in subsequent parts, the number of filters is double the number in the previous part. Since Transformers have 1D vector sequences as input, the feature maps are flattened. The Transformer in the proposed architecture consists of 12 Transformer encoders. Each Transformer encoder mainly consists of two sub-blocks - multi-head self-attention (MSA) and the feed-forward network (FFN), which are connected sequentially. Layer normal (LN) and residual connections are employed before and after two sub-blocks. In order to ensure the consistency of the encoders, the size of the input is the same as the size of the output. The output of the previous encoder is given as input to the next encoder. In the final layer, the embeddings are fed into softmax classifiers to detect significant stenosis.

**RDFNet** approach proposed by Jiang et al. [290] basically incorporates the Transformer mechanism in order to better extract the complex tooth decay features. The incorporation of the

Transformer has improved the detection accuracy. The main three modules of the network are the backbone, neck, and prediction modules. The backbone module is mainly used to extract the features from caries images. In the backbone module, the focus operation is a slicing operation that could easily replace the convolution operation and reduce the loss of feature information. The C3Modified layer is a convolution module activated by the FReLU function, which extracts complex visual-spatial information of the caries images. SPP [302] module has a spatial pyramid structure that could expand the perceptual field, fusing the local and global features and enhancing the feature maps. After the SPP structure, RDFNet appends an improved Transformer-encoder module to improve the feature extraction capability. The main functionality of the neck module is to fuse the feature maps of different sizes and extract high-level semantic structures. This module mainly uses the structure of the feature pyramid network (FPN) proposed in [303], and path aggregation network (PAN) proposed in [304]. The FPN approach is employed in a top-down fashion, and PAN is performed in a bottom-up fashion to generate the feature pyramids. In order to prevent information loss, feature fusion is performed using both bottom-up and top-down approaches. An improved C3Modified convolutional module is adopted into the neck module to better extract the semantic features of caries images. The high-level features generated by the neck module are used by the prediction module, which in turn is used to classify and regress the location and class of the objects. To overcome the problems of the single-stage detection method, which has quite a low detection accuracy, it mainly has three detection heads for detecting large, medium, and small objects. As Transformers have proved to have strong feature extraction capability, in order to extract complex features, they utilized the Transformer model. To better extract the features, three Transformer encoders were stacked together. To simplify the model, the authors removed the original normalization layer from the Transformer encoder. In order to extract the deep features, the feature map was fed into this structure. For each head, the attention values were calculated independently and later concatenated.

Wagner et al. [291] proposed a novel hybrid cell detection approach **(CellCentroidFormer)** in microscopic images that combines the advantages of vision Transformers (ViTs) and convolutional neural networks (CNNs). Authors show that the combined use of convolutional and Transformer layers is advantageous as the convolutional layers can focus on the local information (cell centroids), and the Transformer layers can focus on the global information ( overall shapes of a cell). The proposed centroid-based approach represents the cells as ellipses and is trainable in an end-to-end fashion. Four different 2D microscopic datasets were used for experimental evaluations, and the results outperformed the fully convolutional architecture-based methods. Figure 35 shows the architecture. The encoder is then folded into a 3D tensor, which is afterward concatenated with the input tensor. The MobileViT block is a lightweight alternative to the actual encoder-decoder approach using a Transformer [14]. Due to the multi-head self-attention layers, the MobileViT block causes a much higher computational com-
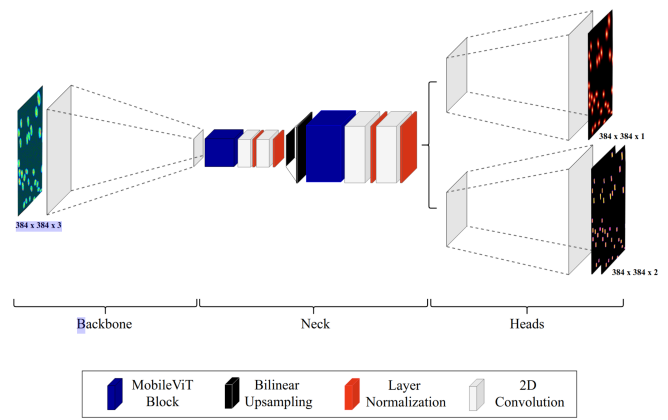


Figure 35: Proposed architecture of CellCentroidFormer model [291].

plexity than convolutional layers. To not increase the computational complexity excessively, the MobileViT blocks are combined in the neck part of the proposed model. Layer normalization is added for regularization and to allow higher learning rates. The backbone module of the proposed model is the EfficientNetV2S [305] CNN model. This block mainly consists of six high-level blocks, out of which five blocks are used to extract image features. To use the advantage of transfer learning, the backbone module is initialized with weights learned from training on ImageNet. This, in turn, reduces the amount of required training data. The EfficientNetV2S [305] CNN models are generally optimized for a fixed input size. Therefore the input images need to be resized to this input size. The cells are represented mainly by the centroid, width, and height parameters. Mainly, two fully convolutional heads are used to predict these cell parameters in the paper. These heads contain 2D convolution, batch normalization, and bilinear upsampling layers. More MobileViT blocks are not used as it will increase the computational complexity. Later convolutional layers have a bigger receptive field which helps in capturing the global information [306] effectively. The first convolutional head predicts a heatmap for detecting the cell centroids, and the second head is used for predicting the cell dimensions. The output dimensions of this model are $384 \times 384$. The authors use one decoder of the Dual U-Net to predict the centroid heatmap, and the second branch predicts the dimensions of the detected cells. The shapes of the cells are focused on by the Transformer layers in the network.

Nguyen et al. [292] present a novel self-supervised learning (SSL) framework, **Joint-2D3D-SSL**, designed to address the scarcity of labeled training data for joint learning on 2D and 3D medical data modalities. The framework constructs an SSL task based on a 2D contrastive clustering problem for distinct classes using 2D images or 2D slices extracted from 3D volumes. In particular, the framework leverages 3D volumes by computing feature embeddings for each slice and then constructing a comprehensive feature representation through a deformable self-attention mechanism. This approach enables the capture of correlations among local slices, resulting in a holistic understanding of the data. The global embedded features derived from this transformer are subsequently utilized to define

an agreement clustering for 3D volumes and a masked encoding feature prediction. By employing this methodology, the framework enables the learning of feature extractors at both the local and global levels, ensuring consistency and enhancing performance in downstream tasks. Experimental results across various downstream tasks, including abnormal chest X-ray detection, lung nodule classification, 3D brain segmentation, and 3D heart structures segmentation, illustrate the efficacy of this joint 2D and 3D SSL approach, surpassing plain 2D and 3D SSL approaches, as well as improving upon SOTA baselines. Additionally, this method overcomes limitations associated with fine-tuning pre-trained models with different dimensionality, providing versatile pre-trained weights suitable for both 2D and 3D applications.

### 7.2. Head

Detection Transformers based on Transformer encoder-decoder architecture require a large amount of training data to deliver the highest performance. However, this is not feasible in the medical domain, where access to labeled data is limited. To address this problem, for the detection of 3D anatomical structures from the human body, Wittmann et al. [296] proposed a detection Transformer network with a **focused decoder**. This network considers the relative position of the anatomical structures and thus requires less training data. The focused decoder uses an anatomical region atlas to deploy query anchors to focus on the relevant anatomical structures. The proposed network omits the Transformer encoder network and consists of only Transformer decoder blocks. The authors show that in 3D datasets, avoiding the encoder can reduce the complexity of modeling relations with a self-attention module.

The model architecture contains a backbone network for feature extraction, a focus decoder network for providing well-defined detection results, a classification network to predict the classes, and a bounding box regression network to output the best possible bounding box. The feature extraction backbone network is a feature pyramid network (FPN) inspired by the RetinaNet [307]. Features from the second layer (P2) are flattened before being given as input to the focus decoder. A specific anatomical region atlas [308] containing regions of interest (RoI) is determined for each dataset. Then to each RoI, uniformly spaced query anchors are placed, and a dedicated object query is assigned to each. Such an object query will restrict the focus decoder network to predict solely within their respective RoI.

The focused decoder network contains a self-attention module, a focused cross-attention module, and a feedforward network (FFN). The self-attention module encodes strong positional inter-dependencies among object queries. The focused cross-attention module matches the input sequence to object queries to regulate the individual feature map for prediction via attention. The FFN network then enables richer feature representation. Also, residual skip connections and normalizations are used to increase gradient flow. The classification network consists of a single fully-connected layer, and the bounding box regression network consists of three layers. The bounding box predictions are combined with query anchors to get the bounding box together with class-specific confidence scores. The net-

work is trained to predict 27 candidates predictions per class. Dynamic labeling with the help of generalized intersection over union (GIoU) is created during training to get 27 predictions. During inference, the prediction with the highest confidence score indicates the best candidate. The model is trained end-to-end with the above GIoU loss, binary cross-entropy loss for the classification network, and L1 loss for the bounding box predictions.

### 7.3. Neck

Detection methods using region-based approaches need to generate anchor boxes to encode their prior knowledge and use a non-maximum suppression to filter the resulting bounding boxes after prediction. These pre-and post-processing steps remarkably reduce the detection performance. To bypass these surrogate tasks, Carion et al. [187] proposed Detection Transformer (DETR), which views the object detection task as a direct set prediction problem using an encoder-decoder architecture using Transformers. The self-attention mechanism of the Transformers, which explicitly models all pairwise interactions between elements in a sequence, helps to predict the set of detections with absolute prediction boxes directly from the image rather than using an anchor. For the end-to-end detection of polyp lesions, Shen et al. [293] proposed a convolution in Transformer **(COTR)** network based on the DETR model. COTR consists of 4 main layers: 1) a CNN backbone network used for extracting features, 2) Transformer encoder layers embedded with convolutional layers used for feature encoding and reconstruction, 3) Transformer decoder layers used for object querying, and 4) a feed-forward network used for detecting prediction. Embedding convolutional layers into the Transformer encoder layer leads to convergence acceleration compared to the slow convergence of the DETR model.

The CNN backbone uses a pre-trained model with ResNet18 [64] architecture for feature extraction. This layer converts input medical images to a high-level feature map. The authors then use a $1 \times 1$ convolution to reduce the channel dimensions. In the Transformer encoder layers, they used six convolution-in-Transformer encoders to collapse this spatial structure into a sequence. Then they use a convolution layer to reconstruct the sequential layer back to the spatial one. In the encoder layer, each Transformer has a standard architecture with a multi-head self-attention module and a feed-forward network. To the input of each attention layer, a positional embedding [14] is also introduced. In the Transformer decoder layers, they used six decoders which follow the standard architecture of the Transformer except that it also decodes object queries in parallel. Each object query will correspond to a particular object in the image. The decoders take these object queries with position embeddings as well as output embeddings from the encoder network and convert them into decoded embeddings. Then they used a feed-forward network with two fully connected layers for converting these decoded embeddings into object predictions. The first fully connected layer is a box regression layer to predict object location, and the second one is a box-classification layer to predict object scores. Therefore, the object queries are independently decoded into box coordinates

Table 14: An overview of the reviewed Transformer-based medical image detection approaches.

| Method | Modality | Organ | Type | Pre-trained Module: Type | Datasets | Metrics | Year |
|---|---|---|---|---|---|---|---|
| **Backbone** | | | | | | | |
| TR-Net [289] | MPR | Heart (Coronary Artery) | 3D | Supervised | Private dataset | Accuracy, Sensitivity, Specificity, PPV, NPV, F1-score | 2021 |
| RDFNet [290] | Dental Caries | Teeth | 2D | Supervised | Private dataset | Precision, Recall, mAP@0:5 | 2021 |
| CellCentroidFormer [291] | Microscopy | Cells | 2D | Supervised | [1] Fluo-N2DL-HeLa (HeLa) [309], [2] Fluo-N2DH-SIM+ (SIM+) [309], [3] Fluo-N2DH-GOWT1 (GOWT1) [309], [4] PhC-C2DH-U373 (U373) [309]. | Mean-IoU, SSIM | 2022 |
| Joint-2D3D-SSL [292] | X-ray | Chest | 2D 3D | Self-supervised | VinDr-CXR [110] | mAP@0.5 | 2022 |
| **Head** | | | | | | | |
| Focused decoder [296] | CT | Multi-organ | 3D | Semi-Supervised | [1] VISCERAL anatomy benchmark [310], [2] AMOS22 challenge [311]. | mAP, AP50, AP75 | 2023 |
| **Neck** | | | | | | | |
| COTR [187] | Colonoscopy | Colon | 2D | Supervised | [1] CVC-ClinicDB [312], [2] ETIS-LARIB [313], [3] CVC-ColonDB [314]. | Precision, Sensitivity, F1-score | 2021 |
| CT-CAD [294] | X-ray | Chest | 2D | Supervised | [1] Vinbig Chest X-Ray dataset [315] [2] ChestXDet-10 dataset [316] | AP50 | 2021 |
| Spine-Transformer [295] | CT | Vertebra | 3D | Supervised | [1] VerSe 2019 [317], [2] MICCAI-CSI 2014 [318], [3] Private dataset | Id-Rate, L-Error | 2021 |

and classes by the feed-forward network, which results in final predictions, including object and no object (background) predictions. This model transforms the object detection problem into a direct set prediction problem by training end-to-end by calculating bipartite matching loss (Hungarian algorithm) between predictions and ground truth for each query. If the number of queries exceeds the number of objects in the image, the remaining boxes are annotated as no object class. Thus, the model is trained to predict output for each query as an object or no object detection. For the class prediction, they used negative log-likelihood loss, and for the bounding box localization, they used an L1 loss with generalized intersection over union (GIOU) [323] loss. The experiments demonstrated that the proposed model achieved comparable performance against state-of-the-art methods.

Many deep learning detection methods lack using context-relevant information for improved accuracy, and they also generally suffer from slower convergence issues and high computational costs. The proposed **CT-CAD** [294], context-aware Transformers for end-to-end chest abnormality detection, address these problems. The model consists of two main modules: 1) a context-aware feature extractor module for enhancing the features, and 2) a deformable Transformer detector module for detection prediction and to accelerate the convergence speed. The context-aware feature extractor network uses a ResNet50 backbone, dilated context encoding (DCE) blocks, and positional encoding structure. The deformable Transformer detector contains a Transformer encoder-decoder architecture and a feed-forward network. The proposed design of the context-aware feature extractor is inspired by the feature fusion scheme from DetectoRS [324] which is based on the Feature Pyramid Networks (FPN) [325]. The feature fusion scheme iteratively

enhances the features of the FPN to powerful feature representations. Likewise, the DCE blocks enhance the features extracted from the ResNet50 backbone by expanding the receptive fields to fuse multiscale context information using dilated convolution filters of different sizes. This powerful feature map benefits in detecting objects across various scales. Inspired by YOLOF [326] the DCE block uses dilated convolution and skip connections to achieve a larger receptive field and acquire more local context information. Finally, all the features from different DCE blocks computed at different scales are summed up to get the feature map for the output.

The proposed design of the deformable Transformer detector contains single-scale and multi-head attention properties. The deformable attention block attends to a small set of key sampling points, thus allowing the Transformer to focus on the feature space and accelerate the convergence. The authors used six encoder and decoder layers with positional encoding to obtain the decoder outputs. The outputs from the decoder are the number of abnormalities detected and the dimension of the decoder layers. Finally, a feed-forward network is used to output the category classification and location regression results. The model is trained end-to-end with a combination of bounding box loss and classification (cross-entropy) loss. The authors adopted GIoU [327] to balance the loss between large and small object bounding boxes.

The attention module in the detection Transformers computes similarity scores between elements of each input data to identify complex dependencies within these data. Calculating similarities of all possible positional pairs in the input data scales quadratically with the number of positions and thus becomes computationally very expensive. For this reason, the Transformer-based object detection model from 3D images

Table 15: A brief description of the reviewed Transformer-based medical image detection models. The unreported number of parameters indicates that the value was not mentioned in the paper.

| Method | # Params | Contributions | Highlights |
| --- | --- | --- | --- |
| **Backbone** | | | |
| TR-Net [289] | - | • This work is the first attempt to detect coronary artery stenosis more accurately by employing Transformers.<br>• To detect significant stenosis, local and global features are effectively integrated into this approach, which has resulted in more accurate results. | • While detecting significant stenosis, the TR-Net architecture is capable of combining the information of local areas near stenoses and the global information of coronary artery branches.<br>• Compared to state-of-the-art methods, the TR-Net model has better results on multiple indicators.<br>• The shallow CNN layer prevents the overfitting of semantic information and improves the overall efficiency.<br>• The gain in performance comes with a trade-off in the number of parameters, which affects the computational complexity. |
| RDFNet [290] | - | • An image dataset of caries is created, which is annotated by professional dentists.<br>• For better extraction of the complex features of dental caries, the Transformer mechanism is incorporated.<br>• In order to increase the inference speed significantly, the FReLU activation function is adopted. | • Compared with existing approaches, the accuracy and speed of caries detection are better.<br>• Method is applicable to portable devices.<br>• The method does not work really well when the illumination of the oral image is insufficient.<br>• Even though detection accuracy and speed are improved compared to the original approach, the detection speed is not the fastest. |
| CellCentroidFormer [291] | 11.5M | • A novel deep learning approach that combines the self-attention of Transformers and the convolution operation of convolutional neural networks is proposed.<br>• A centroid-based cell detection method, denoting the cells as ellipses is proposed. | • Pseudocoloring in combination with pre-trained backbones shows improved cell detection performance.<br>• The model outperforms other state-of-the-art fully convolutional one-stage detectors on four microscopy datasets, despite having a lower number of parameters.<br>• Larger output strides worsen the performance. |
| Joint-2D3D-SSL [292] | 31.16M | • Introduces a self-supervised learning framework capable of leveraging both 2D and 3D data for downstream applications.<br>• Proposes a deformable self-Attention mechanism that captures flexible correlations between 2D slices, leading to powerful global feature representations.<br>• Extends the SSL tasks by proposing a novel 3D agreement clustering and masking embedding prediction | • Overcomes limitations in fine-tuning pre-trained models with different dimensionality.<br>• Produces versatile pre-trained weights for both 2D and 3D applications. |
| **Head** | | | |
| Focused Decoder [296] | VISCERAL Dataset - 41.8M<br>AMOS22 Dataset - 42.6M | • First detection Transformer model for 3D anatomical structure detection.<br>• Introduced a focused decoder to focus the predictions on RoI. | • Better results compared to existing detection models using a Transformer network like DETR [187] and deformable DETR [16].<br>• Comparable results to the RetinaNet [307].<br>• Varying anatomical fields of view (FoVs) can affect the robustness of the model. |
| **Neck** | | | |
| COTR [162] | - | • Proposed a convolution layer embedded into the Transformer encoder for better feature reconstruction and faster convergence compared to DETR. | • COTR has comparable results with state-of-the-art methods like Mask R-CNN [319] and MDeNet-plus [320].<br>• This approach produces low confidence for a particular type of lesion. |
| CT-CAD [294] | - | • Proposed a context-aware feature extractor, which enhances the receptive fields to encode multi-scale context-relevant information.<br>• Proposed a deformable Transformer detector that attends to a small set of key sampling locations and then the Transformers can focus to feature subspace and accelerate the convergence speed. | • CT-CAD outperforms the existing methods in Cascade R-CNN [321], YoLo [322], and DETR [16].<br>• CT-CAD is capable to detect hard cases, such as nodules that are ignored by Faster R-CNN.<br>• Compared to the ChestXDet-10 dataset, this model has a lower performance on the Vinbig Chest X-Ray dataset which has higher categories of abnormalities with more complex patterns. |
| Spine-Transformers [295] | - | • Proposed a 3D object detection model based on the Transformer's architecture.<br>• Proposed a one-to-one set global loss that enforces unique prediction and preserves the sequential order of vertebrae.<br>• Proposed a Sphere-based bounding box to enforce rotational invariance. | • Obtained better results for all datasets compared to state-of-the-art methods.<br>• The model has a higher Id-Rate on both the datasets, but a higher L-Error compared to the benchmark by [317]. |

has never been applied. Tao et al. [295] proposed a novel Transformer-based 3D object detection model as a one-to-one set prediction problem for the automatic detection of vertebrae in arbitrary Field-Of-View (FOV) scans, called the **Spine-Transformers**. Here the authors used a one-to-one set-based global loss that compels a unique prediction for preserving the sequential order of different levels of vertebrae and eliminated bipartite matching between ground truth and prediction. The main modules of the Spine-Transformer are (1) a backbone network to extract features, (2) a light-weighted Transformer encoder-decoder network using positional embeddings and a skip connection, and (3) two feed-forward networks for detection prediction. The authors used a ResNet50 [64] architecture without the last SoftMax layer as the backbone network to extract high-level features. These features are passed through a $1 \times 1 \times 1$ convolutional layer to reduce the channel dimensions and then flattened to get a feature sequence to feed as the input for the Transformer network. The light-weighted Trans-

former encoder-decoder network contains only a two-layer encoder and two-layer decoder to balance between feature resolution and memory constraint. In both the encoder and decoder layers of the network, learnable positional embeddings are used. The authors found that using a skip connection across the Transformer encoder-decoder network will help in the propagation of context and gradient information during training and thus improves performance. The two feed-forward networks are then used to predict the existence of the objects and regress their coordinates. The authors also proposed a sphere-based bounding box detector to replace the rectangular-based bounding box to introduce rotational invariance called InSphere detector. The Spine-Transformer is trained end-to-end with fixed-size patch images to predict all the vertebrae objects in parallel by forcing one-to-one matching. Binary cross-entropy loss is used as classification loss, and to enforce the order of the predicted vertebrae objects, an edge loss is introduced, which is an L1 distance loss introduced between the centers of the top

and bottom neighborhood vertebrae objects. For better localization accuracy of the bounding sphere detection, the authors used generalized inception-over-union (GIoU) [323] loss. The results of this model showed superior results to all the state-of-the-art methods. The authors also claim that by using a 3D CNN-based landmark regression [328], the localization accuracy can be further improved.

### 7.4. Discussion and Conclusion

In this chapter, several well-known Transformer architectures are analyzed to address the automatic detection challenge. Based on the Transformer model contribution to the network structure, we grouped the set of literature work into the backbone, neck, or head strategies and for each category, we provided sample works. We also described the details regarding self-supervised learning in [292]. In this respect, the core idea behind each network design along with the pros and cons of the strategies are highlighted in the summary tables. Vision Transformers have been shown to make more accurate diagnoses compared to traditional methods of analyzing medical images. These deep learning models can be trained on large datasets, such as ImageNet, and fine-tuned on medical image datasets to improve their performance in detecting abnormalities in X-rays, CT scans, and MRIs. By incorporating information from multiple modalities, Transformers can further enhance their ability to identify and detect rare or subtle abnormalities in medical images. Many medical images are often taken over time, and incorporating temporal information into the model can improve its performance. For example, the model can be designed to take into account the temporal evolution of diseases or conditions. Overall, Transformers have demonstrated their capabilities to significantly improve the accuracy and efficiency of medical image analysis, leading to advances in healthcare.

## 8. Medical Image Registration

Medical image registration is the task of transforming a set of two or more images of an organ or a biological process taken with different poses, time stamps, or modalities (e.g., CT and MRI) into a geometrically aligned and spatially corresponding image that can be utilized for medical analysis. The transformation can be discovered by solving an optimization problem that maximizes the similarity between the images to be registered [329]. A pair-wise registration of two MRI brain scans is shown in Figure 36 for illustration.

Despite remarkable advancements in the quality of medical imaging techniques that aid professionals in better visualization and analysis of image data, a prominent challenge prevails in developing a system capable of effective integration of visual data that captures useful information from original images with high precision. Most registration procedures take into account the whole image as input by utilizing global information for spatial transformation, which leads to inefficient and slow integration of data. Furthermore, the collection process of medical images for training is slow and toilsome, performance degrades
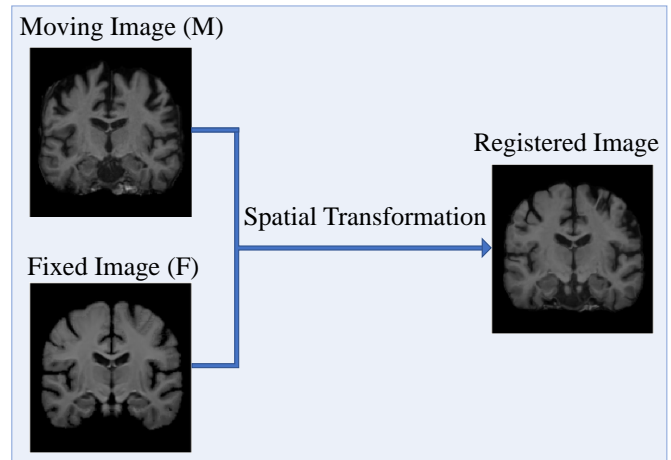


Figure 36: An example of pair-wise medical image registration. The goal of image registration is to geometrically align the moving image with the target or fixed image by performing the spatial transformation.

due to the presence of outliers, and local maxima entail negative effects on performance during optimization [330, 331]. The emergence of deep learning methods alleviated these problems by automatic extraction of features utilizing convolutional neural networks (CNN), optimizing a global function, and improving registration accuracy. For instance, Balakrishnan et al. [332] utilized a CNN to achieve unsupervised deformable registration by treating it as a parametric function to be optimized during training. Furthermore, Chen et al. [333] presented an unsupervised CNN-based registration algorithm to produce anthropomorphic phantoms. However, there are still limitations in capturing long-range spatial correspondence in CNN-based frameworks [334, 39].

Fueled by the strong ability of Transformers to model long-range dependencies and detect global information [335, 21, 336], they have gained the attention of researchers in the medical image registration domain in recent years. In this section, we review Transformer-based methods in medical image registration that ameliorate the aforementioned shortcomings of previous systems by utilizing the self-attention mechanism. We have organized the relevant approaches based on their type of registration:

(a) Deformable registration, which employs an optimization algorithm to tune the transformation model, is a way that maximizes the similarity measure function for the images of interest [337];

(b) Rigid registration, which achieves correspondence by maintaining the relative distance between each pair of points between the patient's anatomy images [337].

(c) Affine registration, which contains the same operations as rigid registration plus non-isometric scaling.

### 8.1. Deformable Registration

Most existing Transformer-based algorithms focus on deformable transformation to perform medical image registration. **Vit-V-Net** [338] is the earliest work that incorporates
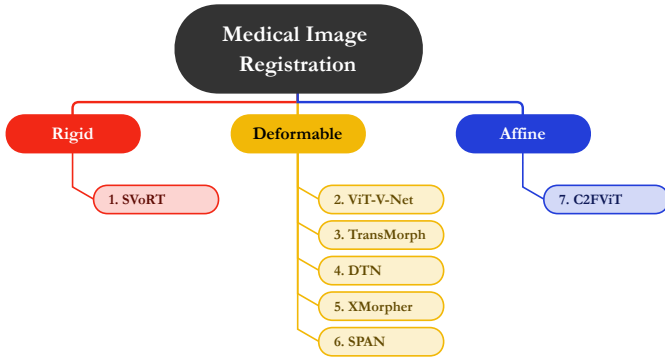
Figure 37: Taxonomy of Transformer-based image registration based on their transformation type. We use the prefix numbers in the figure in ascending order and reference the corresponding paper as follows: 1. [336], 2. [338], 3. [339], 4. [340], 5. [341], 6. [342], 7. [343].
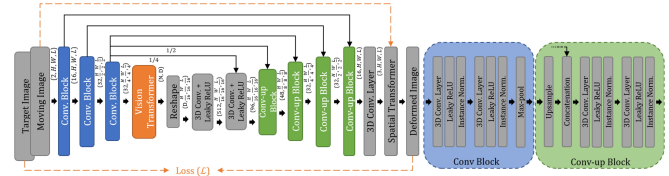


Figure 38: Overview of ViT-V-Net. Multiple convolution blocks encode images into high-level features, which the Vit block splits into patches. These patches are then mapped to D-dimensional patch embeddings that get integrated with learnable positional encodings to retain positional information. Next, these patches are passed into the Transformer encoder block, followed by multiple skip connections to retain localization information, and decoded using a V-Net style decoder. Using the network's final output, a spatial Transformer warps the moving image. Figure taken from [338].

Transformers to perform medical image registration in a self-supervised fashion. It is inspired by the integration of vision Transformer-based segmentation methods with convolutional neural networks to enhance the localization information recovered from the images. Unlike previous research that employed 2D images for spatial correspondence, Vit-V-net stepped towards utilizing ViT [15] as the first study for volumetric medical image registration (i.e., 3D image registration). As illustrated in Figure 38, the images are first encoded into high-level feature representations by implementing multiple convolution blocks; then, these features get split into *P* patches in the ViT block. Next, the patches are mapped to a D-dimensional embedding space to provide patch embeddings, which are then integrated with learnable positional encodings to retain positional information. Next, these patches are passed into the encoder block of the Transformer, followed by multiple skip connections to retain localization information, and then decoded employing a V-Net style decoder [344]. Finally, a spatial Transformer [345] warps the moving image by utilizing the final output of the network. TransMorph [339] extended ViT-V-Net and proposed a hybrid Transformer ConvNet framework that utilizes the Swin Transformer [58] as the encoder and a ConvNet as the decoder to provide a dense displacement field. Like ViT-V-Net, it employed long skip connections to retain the flow of localization information that may enhance registration accuracy. The output of the network, which is a nonlinear warping function, gets applied to the moving image with the deformation field utilizing the spatial transformation function proposed in [345]. An affine transformation Transformer network is incorporated to align the moving image with the fixed image before feeding it to the deformable registration network. This work also proposed two variants of TransMorph: diffeomorphic TransMorph (TransMorph-diff) to facilitate topology-preserving deformations and Bayesian TransMorph (TransMorph-Bayes) to promote a well-calibrated registration uncertainty estimate.

Likewise, Zhang et al. [340] introduced the dual Transformer network (**DTN**) framework to perform diffeomorphic registration. It is composed of a CNN-based 3D U-Net encoder [328] for the embedding of separate and concatenated volumetric images and a dual Transformer to capture the cross-volume dependencies. One of the Transformers is responsible for modeling the inter- and intra-image dependencies, and the other one handles the modeling of the global dependencies by employing the self-attention mechanism. The concatenation of the generated features from these Transformers results in enhanced feature embeddings, which are utilized by the CNN-based decoder to provide a diffeomorphic deformation field. The evaluation of the framework was conducted on the brain MRI scans of the OASIS dataset [346], which substantiates their improvements in diffeomorphic registration compared to the existing deep-learning-based approaches.

Furthermore, **XMorpher** [341] put emphasis on the significance of backbone architectures in feature extraction and match of pair-wise images, and proposed a novel full Transformer network as the backbone, which consists of two parallel U-Net structures [28] as the sub-networks with their convolutions replaced by the introduced Cross Attention Transformer for feature extraction of moving and fixed images, and cross-attention-based fusion modules that utilize these features for generating the feature representation of moving-fixed correspondence and fine-grained multi-level semantic information that contributes to a fine registration.

The latest developments in self-supervised representation learning have primarily aimed at eliminating inductive biases in training processes. Nonetheless, these biases can still serve a purpose in scenarios with scarce data or when they offer further understanding of the underlying data distribution. To mitigate this problem, [342] introduces spatial prior attention (**SPAN**), a framework that leverages the consistent spatial and semantic structure found in unlabeled image datasets to guide the attention mechanism of vision Transformers. **SPAN** integrates domain expertise to enhance the effectiveness and interpretability of self-supervised pretraining specifically for medical images using image registration and alignment methodologies. Specifically, they exploit image registration to align the attention maps with an inductive bias corresponding to a salient region so only a single representative sample is required. When utilizing deformable image registration templates, **SPAN** achieves the highest mAUC score, followed by the predicted global templates and the triangular spatial heuristic resulting in improved performance in downstream tasks.

Table 16: An overview of the reviewed Transformer-based medical image registration approaches.

| Method | Modality | Organ | Type | Datasets | Metrics | Year |
|---|---|---|---|---|---|---|
| **Deformable** | | | | | | |
| ViT-V-Net [338] | MRI | Brain | 3D | Private Dataset | Dice | 2021 |
| TransMorph [339] | MRI CT | Brain Chest-Abdomen-Pelvis region | 3D | [1] IXI [265] [2] T1-weighted brain MRI scans from Johns Hopkins University [3] Chest-Abdomen-Pelvis CT [347] | Dice % of $|J_\Phi| \leq 0$ SSIM | 2022 |
| DTN [340] | MRI | Brain | 3D | OASIS [346] | Dice $|J_\Phi| \leq 0$ | 2021 |
| XMorpher [341] | CT MRI | Heart | 3D | [1] MM-WHS 2017 [348] [2] ASOCA [349] | Dice % of $|J_\Phi| \leq 0$ | 2022 |
| SPAN [342] | CT MRI | Lung | 3D | CheXpert [75] JSRT [350] | mAUC | 2022 |
| **Affine** | | | | | | |
| C2FViT [343] | MRI | Brain | 3D | [1] OASIS [346] [2] LPBA [351] | Dice Hausdorff distance | 2022 |
| **Rigid** | | | | | | |
| SVoRT [336] | MRI | Brain | 3D | FeTA [352] | PSNR SSIM | 2022 |

Table 17: A brief description of the reviewed Transformer-based medical image registration techniques.

| Method | Contributions | Highlights |
|---|---|---|
| **Deformable** | | |
| ViT-V-Net [338] | • Contributed to the medical image registration domain as the first work to exploit ViTs to develop a volumetric (3D) registration. • Integrated Transformers with CNNs to build a hybrid architecture for self-supervised brain MRI registration. | • Employed a hybrid architecture to incorporate long-range and local information in the registration process. • Attempted to preserve the localization data with the help of long skip connections between the encoder and decoder stages. |
| TransMorph [339] | • Proposed a Transformer-based unsupervised registration approach for affine and deformable objectives. • Conducted experiments on two brain MRI datasets and in a phantom-to-CT registration task to demonstrate their superior performance compared to traditional approaches. | • They additionally proposed two distinguishable versions of their model: a diffeomorphic variant to facilitate the topology-preserving deformations and a Bayesian variant to promote a well-calibrated registration uncertainty estimate. • Studied the effect of receptive fields by comparing TransMorph with CNNs and addressed that while the receptive field of ConvNets only increases with the layer depth, their presented model takes into account the whole image at each due to the self-attention mechanism. |
| DTN [340] | • Proposed a dual Transformer architecture to capture semantic correspondence of anatomical structures. • The suggested DTN demonstrated remarkable results in diffeomorphic registration and atlas-based segmentation of multi-class anatomical structures. | • The Dual Transformer is capable of reducing the negative Jacobian determinant while preserving the atlas-based registration quality. • The qualitative and quantitative analysis of their method on the OASIS dataset indicates that diffeomorphic registration fields are effective. |
| XMorpher [341] | • Devised a deformable registration system consisting of dual parallel feature extraction networks which facilitate the association of representative features between moving and fixed images. • Proposed cross-attention Transformer that establishes spatial correspondences through computation of bilateral information in the attention mechanism. | • Promotes registration superiority by presenting fine-grained visual results in terms of boundary smoothness, adjacent regions' resolution quality, and deformation grid polishness. • Demonstrated the model's great diagnostic potential by conducting experiments with different training regimes. |
| SPAN [342] | • Developed a self-supervised transformer-based registration method that leverages consistent spatial and semantic structure in unlabeled image datasets to guide the attention mechanism of vision Transformers. | • Integrates domain expertise and uses image registration and alignment methodologies to enhance the effectiveness and interpretability of self-supervised pretraining, particularly for medical images. |
| **Affine** | | |
| C2FViT [343] | • Presented a method in order to learn the global affine registration by taking advantage of the strong long-range dependency recognition and locality of the hybrid Transformer and the multi-resolution strategy. | • The suggested training framework can be extended to a number of parametric-based registration approaches by removing or scaling the geometrical transformation matrices. |
| **Rigid** | | |
| SVoRT [336] | • Devised an approach for the task of Volumetric reconstruction of fetal brains based on Transformer architectures. • Employed a Transformer network trained on artificially sampled 2D MR slices that estimate the underlying 3D volume from the input slices to more accurately predict transformation. | • Experimental procedures on the FeTA dataset [352] represented the model's ability in high-quality volumetric reconstruction. • The volumetric reconstruction associated with the transformations of the proposed method displays higher visual quality. |
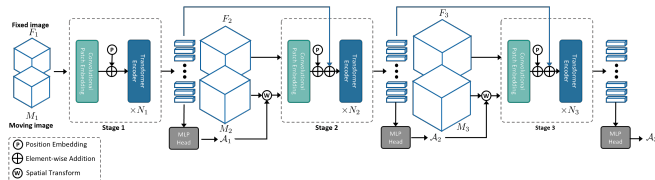


Figure 39: The model has $L$ stages with convolutional patch embedding layers and $N$ Transformer encoder blocks to learn the optimal affine registration matrix. In each stage, fixed and moving images are downsampled and concatenated, then passed to the convolutional patch embedding layer to produce image patch embeddings. The Transformer then produces the input feature embedding from the embeddings [343].

### 8.2. Affine Registration

To perform affine medical image registration with Transformers, Mok et al. [343] proposed **C2FViT**, a coarse-to-fine vision Transformer that performs affine registration, a geometric transformation that preserves points, straight lines, and planes while registering 3D medical images. Former studies have relied on CNN-based affine registration that focuses on local misalignment or global orientation [353, 354], which limits the modeling of long-range dependencies and hinders high generalizability. C2FVit, as the first work that takes into account

the non-local dependencies between medical images, leverages vision Transformers instead of CNNs for 3D registration. As depicted in Figure 39, the model is split into $L$ stages, each containing a convolutional patch embedding layer and $N_i$ Transformer encoder blocks (i indicates the stage number), intending to learn the optimal affine registration matrix. In each stage, the fixed and moving images are downsampled and concatenated with each other, then the new representation gets passed to the convolutional patch embedding layer to produce image patch embeddings. Next, the Transformer receives the embeddings and produces the feature embedding of the input. Conducted experiments on OASIS [346] and LPBA [351] demonstrated their superior performance compared to existing CNN-based affine registration techniques in terms of registration accuracy, robustness, and generalization ability.

### 8.3. Rigid Registration

**SVoRT** [336] addressed the necessity of slice-to-volume registration before volumetric reconstruction for the task of volumetric fetal brains reconstruction, and employed a Transformer network trained on artificially sampled 2D MR slices that learn to predict slice transformation based on the information gained

from other slices. The model also estimates the underlying 3D volume from the input slices to promote higher accuracy in transformation prediction. The superiority of their proposed method in terms of registration accuracy and reconstruction based on the evaluation of synthetic data and their experiments on real-world MRI scans demonstrated the ability of the model in high-quality volumetric reconstruction.

---

### 8.4. Discussion and Conclusion

According to the research discussed in this section, vision Transformers are prominent tools in image registration tasks due to their training capability on large-scale data, which is made feasible by parallel computing and self-attention mechanisms. Leveraging Transformers to encourage better global dependency identification improves registration in terms of dice scores and Jacobian matrix determinants compared to CNNs.

To mitigate the burden of quadratic complexity when processing images at high resolution and modeling local relationships, reviewed studies usually employ CNNs to provide feature maps or dense displacement fields [338, 339, 340]. C2FViT [343] disregarded convolutional networks and implemented convolutional patch embeddings to promote locality. However, in deformably registering medical content, XMorpher recently demonstrated the power of cross-attention in better capturing spatial relevancy without a CNN implementation [341], and SVoRT purely utilized Transformers to perform rigid registration [336]. Regarding the challenges caused by the data-hungry nature of deep neural networks and the difficulties in gathering annotated training data, studies in medical image registration have also focused on integrating self-supervised learning with transformer models [342].

The notable experimental attempts on brain MRI scan data, such as OASIS [346] and FeTA [352], show the importance of accurate automatic registration for neuroimaging data. One particular work [341] proposed to evaluate their registration on images of cardiac region datasets including MM-WHO-2017 [348] and ASOCA [349]. To further clarify the modality type used in the aforementioned proposed methods, all works conducted their evaluations on 3D or volumetric imaging modalities.

Based on the brief review of Transformer-based medical image registration research, we believe that other regions of interest (ROI) such as neurons, retina, and neck area are worth exploring to facilitate diagnostic operations in different domains with more precise registration models.

We have also specified the architectural type, modality, organ, data size, training paradigm, datasets, metrics, and year for each medical registration technique reviewed in Table 16. Furthermore, Table 17 provides a list of the contributions and highlights of the proposed works. We also elaborate on the use cases of self-supervision in medical image registration.

---

## 9. Medical Report Generation

Medical report generation focuses on producing comprehensive captions and descriptions pivoting on medical images for diagnostic purposes. Designing automatic methods capable of performing this task can alleviate tedious and time-consuming work in producing medical reports and promote medical automation [372]. Recently, advancements in deep learning have brought the attention of researchers to employing an intelligent system capable of understanding the visual content of an image and describing its comprehension in natural language format [373]. Research efforts in improving this area can be employed in medical imaging by implementing systems capable of providing descriptions and captions (i.e., generating medical reports) concerning medical images. These captioning systems usually utilize encoder-decoder models that encode medical images and decode their understandings to provide diagnostic information in a natural language format.

Despite the success of deep learning, limitations including reliability on an immense amount of data, unbalanced data in radiology datasets (e.g., IU X-ray chest X-Ray [374]), and the black box nature of DL models entail challenges in medical report generation [375]. The success of Transformer models in many vision-and-language tasks has drawn the attention of researchers in the medical report generation domain to the employment of this architecture. In this section, we discuss approaches that utilize Transformers to promote effective capture of long-range context dependencies and better report generation. As illustrated in Figure 40, the following is our taxonomy of these systems according to the mechanism by which they produce accurate and reliable clinical reports:

(I) *Reinforcement Learning-based.* The ultimate goal of a medical report generation system is to provide clinically accurate and reliable reports. In reinforcement learning, the MRG system is considered an agent with the objective of maximizing clinical accuracy based on the feedback given by the reward signal, which is directly calculated by the evaluation metric score (e.g., CIDEr [376]).

(II) *Graph-based.* Radiology reports are typically composed of a long finding section with multiple sentences that make report generation a challenging task. Therefore, the inclusion of prior information is beneficial for facilitating the generation of long narratives from visual data. Knowledge graphs, which are powerful models that can capture domain-specific information in a structured manner, can be used to exploit prior information for medical report generation [377, 362, 378].

(III) *Memory-based.* Memory is a resource through which important information is recorded. In designing a proper MRG system, it is crucial to store vital and diagnostic information that can benefit the generation process by incorporating prior knowledge and experience. Hence, configuring a memory mechanism with Transformers as a report generation framework facilitates longer and more coherent text generation by sharing information gained through the process [357, 361].

(IV) *Other Systems.* Systems that introduce different ideas from previous categories to improve clinical accuracy,
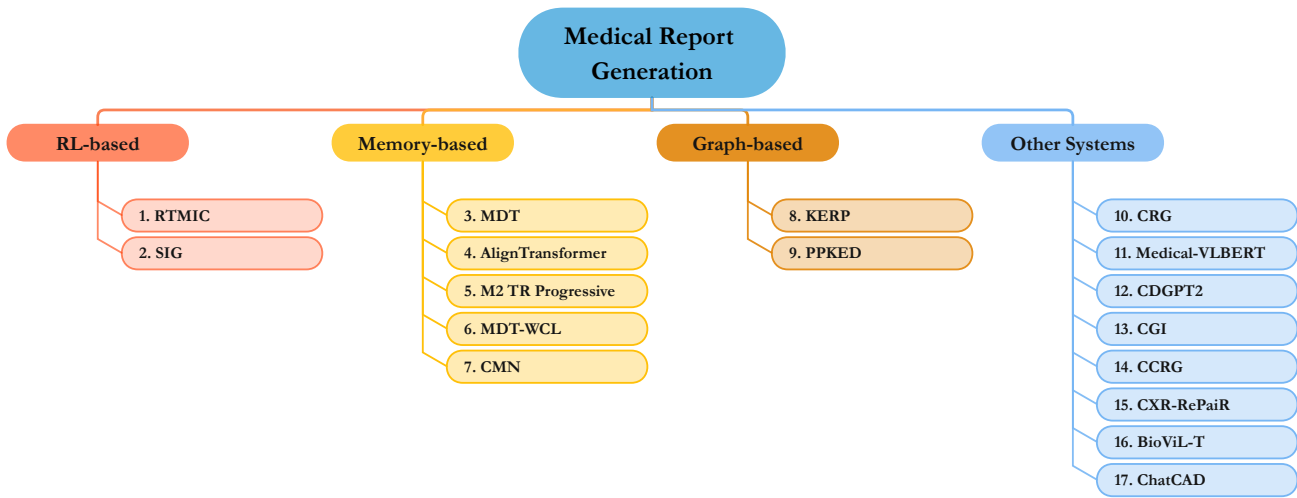
Figure 40: Taxonomy of Transformer-based medical report generation approaches based on the mechanism by which they generate clinical reports. We reference the papers in ascending order corresponding to their prefix number: 1. [355], 2. [356], 3. [357], 4. [358], 5. [359], 6. [360]. 7. [361], 8. [362], 9. [363], 10. [364], 11. [365], 12. [366], 13. [367], 14. [368], 15. [369], 16. [370], 17. [371].

such as curriculum learning, contrastive learning, and alternate learning, belong to this group.

### 9.1. Reinforcement Learning-based Systems

The first work to implement a Transformer architecture for medical report generation is RTMIC [355]. It used the reinforcement learning strategy in training to mitigate the problem of exposure bias prevailing in Seq2Seq models [393]. In their approach, the original images are fed into a DenseNet [379] as the region detector to extract bottom-up visual features. These features are then passed into a visual encoder to generate visual representations from the detected regions, which the captioning detector then utilizes to generate captions for the specified regions. The proposed method was experimented on the IU X-Ray dataset [374] and achieved state-of-the-art results. Integration of RL and Transformers was also applied in surgical instruction generation since the joint understanding of surgical activity along with modeling relations linking visual and textual data is a challenging task. Zhang et al. [356] employed a Transformer-backboned encoder-decoder architecture and applied the self-critical reinforcement learning [394] approach to optimize the CIDEr score [376] as the reward. Their approach surpasses existing models in performance on the DAISI dataset [381] with caption evaluation metrics applied to the model. This work's key difference from others is that their model is proposed to generate instructions instead of descriptions.

### 9.2. Graph-based Systems

In graph-based medical report generation, Li et al. [362] proposed **KERP**, a Graph Transformer implementation to generate robust graph structures from visual features that are extracted by a DenseNet [379] backbone. This approach is composed of three modules: Encode, Retrieve and Paraphrase. First, it constructs an abnormality graph by converting the visual features extracted from the medical images via an encoder module. Next, a sequence of templates is retrieved considering the

detected abnormalities by utilizing a retrieve module. Subsequently, the terms of the produced templates are paraphrased into a report by employing the paraphrase module. The KERP's workflow is illustrated in Figure 41.
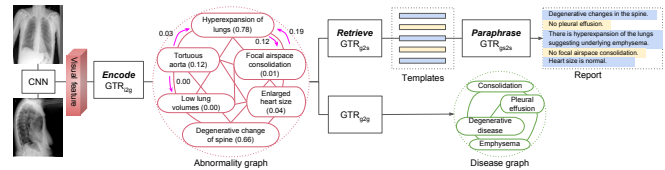


Figure 41: Using an encoder module, KERP creates an abnormality graph from the extracted visual features. Then, a retrieval module retrieves a sequence of templates based on detected abnormalities. Next, the paraphrase module paraphrases the templates' terms into a report [362].

Additionally, Liu et al. [363] addressed the visual and textual data biases and their consequences in generating radiology reports and proposed the **PPKED** framework to alleviate these challenges. Their work introduced three modules to perform report generation: (1) Prior Knowledge Explorer (PrKE), which obtains relevant prior information for the input images; (2) Posterior Knowledge Explorer (PoKE), which extracts the posterior information, including the abnormal regions of the medical image; and (3) Multi-domain Knowledge Distiller (MKD), which distills the obtained information from the previous modules to perform the final report generation. PPKED then formulated the problem by employing the presented modules in the following manner: PoKE first extracts the image features corresponding to the relevant disease topics by taking the visual features extracted by ResNet-152 [64] from the input image and abnormal topic word embeddings as the input. Next, the PrKE module filters the prior knowledge from the introduced prior working experience (a BERT encoder) and prior medical knowledge component that is relevant to the abnormal regions of the input image by utilizing the output of the PoKE module. Next, the MKD module generates the final medical report by using this obtained information, which is implemented based

Table 18: An overview of the reviewed Transformer-based Medical Report Generation approaches.

| Method | Modality | Organ | Type | Visual Backbone | Datasets | Metrics | Year |
|---|---|---|---|---|---|---|---|
| **Reinforcement Learning** | | | | | | | |
| RTMIC [355] | X-ray | Lung | 2D | DenseNet-121 [379] | IU Chest X-ray [374] | BLEU [380]<br>CIDEr [376] | 2019 |
| SIG [356] | Ultrasound<br>Colonoscopy | Multi-organ | 3D | ResNet-101 [64] | DAISI [381] | BLUE [380]<br>Meteor [382]<br>CIDEr [376]<br>ROUGE [383]<br>SPICE [384] | 2021 |
| **Graph** | | | | | | | |
| KERP [362] | X-ray | Lung | 2D | DenseNet-121 [379] | [1] IU Chest X-ray [374]<br>[2] CX-CHR (private dataset) | BLEU [380]<br>CIDEr [376]<br>ROUGE [383] | 2019 |
| PPKED [363] | X-ray | Lung | 2D | ResNet-152 [64] | [1] IU Chest X-ray [374]<br>[2] MIMIC-CXR [385] | BLUE [380]<br>Meteor [382]<br>CIDEr [376]<br>ROUGE [383] | 2021 |
| **Memory** | | | | | | | |
| MDT [357] | X-ray | Lung | 2D | ResNet-121 [64] | [1] IU Chest X-ray [374]<br>[2] MIMIC-CXR [385] | BLEU [380]<br>Meteor [382]<br>ROUGE [383] | 2020 |
| AlignTransformer [358] | X-ray | Lung | 2D | ResNet-50 [64] | [1] IU Chest X-ray [374]<br>[2] MIMIC-CXR [385] | BLEU [380]<br>Meteor [382]<br>ROUGE [383] | 2021 |
| $M^2$ TR. progressive [359] | X-ray | Lung | 2D | DenseNet-121 [379] | [1] IU Chest X-ray [374]<br>[2] MIMIC-CXR [385] | BLEU [380]<br>Meteor [382]<br>ROUGE [383] | 2021 |
| MDT-WCL [360] | X-ray | Lung | 2D | ResNet [64] | [1] MIMIC-ABN [386]<br>[2] MIMIC-CXR [385] | BLUE [380]<br>Meteor [382]<br>ROUGE [383] | 2021 |
| CMN [361] | X-ray | Lung | 2D | ResNet-101 [64] | [1] IU Chest X-ray [374]<br>[2] MIMIC-CXR [385] | BLEU [380]<br>Meteor [382]<br>ROUGE [383] | 2022 |
| **Other** | | | | | | | |
| CRG [364] | X-ray | Lung | 2D | DenseNet-121 [379] | MIMIC-CXR [385] | BLUE [380]<br>Meteor [382]<br>CIDEr [376]<br>ROUGE [383] | 2020 |
| Medical-VLBERT [365] | CT<br>X-ray | Lung | 2D | DenseNet-121 [379] | [1] Chinese Covid-19 CT [387]<br>[2] CX-CHR (private dataset) | BLUE [380]<br>CIDEr [376]<br>ROUGE [383] | 2021 |
| CDGPT2 [366] | X-ray | Lung | 2D | DenseNet-121 [379] | IU chest X-ray [374] | BLUE [380]<br>Meteor [382]<br>CIDEr [376]<br>ROUGE [383] | 2021 |
| CGI [367] | X-ray | Lung | 2D | DenseNet-121 [379] | [1] MIMIC-CXR [385]<br>[2] IU chest X-ray [374] | BLUE [380]<br>Meteor [382]<br>ROUGE [383] | 2021 |
| CGRG [368] | X-ray | Lung | 2D | ResNet-101 [64] | [1] IU Chest X-ray [374]<br>[2] COV-CTR [388] | BLUE [380]<br>Meteor [382]<br>ROUGE [383] | 2021 |
| CXR-RePaiR [369] | X-ray | Lung | 2D | ✗ | MIMIC-CXR [385] | BLUE [380]<br>F1 | 2021 |
| BioViL-T [370] | X-ray | Lung | 2D | ResNet-50 [64] | [1] MIMIC-CXR [385]<br>[2] MS-CXR-T [370] | BLUE [380]<br>ROUGE [383]<br>CHEXBERT [389] | 2023 |
| ChatCAD [371] | X-ray | Lung | 2D | ✗ | MIMIC-CXR [385] | Precision, Recall, F1 | 2023 |

on the decoder part of the Transformers equipped with Adaptive Distilling Attention.

Table 19: A brief summary of the reviewed Transformer-based medical report generation methods.

| Method | Contributions | Highlights |
|---|---|---|
| **Reinforcement Learning** | | |
| RTMIC [355] | • Presented a novel Hierarchical Reinforced Transformer for producing comprehensible, informative medical reports by training through reinforcement learning-based training. <br> • The initial attempt at incorporating Transformers to develop a medical report generation system. | • Utilized reinforcement learning to ameliorate the exposure bias problem. <br> • Enhanced clinical report coherence by employing Transformers to capture long-range dependencies. <br> • The selected metric (CIDEr) as a reward signal is not designed for the medical domain [390]. |
| SIG [356] | • Generated surgical instructions from multiple clinical domains by utilizing a Transformer-based Encoder-decoder architecture. | • The proposed method is able to produce multimodal dependencies, form pixel-wise patterns, and develop textual associations for the masked self-attention decoder. <br> • Utilizing self-critical reinforcement learning to perform optimization increased the performance of surgical instruction generation. <br> • The selected metric (CIDEr) as a reward signal is not designed for the medical domain [390]. |
| **Graph** | | |
| KERP [362] | • Developed an MRG system using a hybrid retrieval-generation technique that unifies standard retrieval-based and recent visual text generation methods. <br> • Introduced Graph Transformer (GTR) as the first research to employ an attention mechanism to convert different data types formulated as a graph. | • Aligning the generated reports with abnormality attention maps by providing location reference facilitates medical diagnosis. <br> • Since KERP is designed based on abnormality detection, it may disregard other valuable information [367]. |
| PPKED [363] | • Proposed a three-module system that mimics the working habits of radiologists by extracting abnormal regions, encoding prior information, and distilling the useful knowledge to generate accurate reports. | • Provides abnormal descriptions and locations to facilitate medical diagnosis. <br> • Capable of extracting relevant information from the explored posterior and prior multi-domain knowledge. <br> • Some mistakes, such as duplicate reports and inaccurate descriptions, are present in the generated reports. [391] |
| **Memory** | | |
| MDT [357] | • Introduced Memory-Driven Transformer for radiology report generation <br> • Developed a relational memory to retain essential knowledge gathered through the previous generations. | • To facilitate medical diagnosis, visual-textual attention mappings were incorporated to capture correspondence with essential medical terms. <br> • Dataset imbalance with dominating normal findings hinders the model's generalizability. |
| AlignTransformer [358] | • Introduced an MRG framework that mitigates the problem of data bias by hierarchically aligning visual abnormality regions and illness tags in an iterative fashion. | • Conducted experiments on MIMIC-CXR and IU-Xray datasets and demonstrated the capability of the model in ameliorating the data bias problem |
| $M^2$ TR. progressive [359] | • Developed a progressive text generation model for medical report generation by incorporating high-level concepts into the process of generation. | • The division of report generation into two steps enhanced the performance in terms of language generation and clinical efficacy metrics. <br> • The progressive generation process increases the false positive rate by including abnormality mentions in negation mode. [390]. |
| MDT-WCL [360] | • Introduced the contrastive learning technique into chest X-ray report generation by proposing a weakly supervised approach that contrasts report samples against each other to better identify abnormal findings. | • Optimization with contrastive loss facilitates generalizability in comparison to contrastive retrieval-based methods. |
| CMN [361] | • Cross-modal memory networks were introduced to improve report generation based on encoder-decoder architectures by incorporating a shared memory to capture multi-modal alignment. | • Capable of properly aligning data from radiological images and texts to aid in the preparation of more precise reports in terms of clinical accuracy. |
| **Other** | | |
| CRG [364] | • Formulated the problem in two steps: (1) a report generation phase incorporating a standard language generation objective to train a Transformer model, and (2) a sampling phase that includes sampling a report from the model and extracting clinical observations from it. | • Transformers' ability to provide more coherent and fluent reports was demonstrated. <br> • Due to the biased nature of the dataset caused by dominant normal findings, the algorithm tends to generate reports that lack essential descriptions of abnormal sections [392]. |
| Medical-VLBERT [365] | • Proposed a framework as the first work that generates medical reports for the COVID-19 CT scans <br> • Devised an alternate learning strategy to minimize the inconsistencies between the visual and textual data. | • Alleviated the shortage of COVID-19 data by employing the transfer learning strategy. <br> • Capable of effective terminology prediction <br> • Overreliance on predetermined terminologies undermines robustness and generalizability. |
| CDGPT2 [366] | • Presented a conditioning mechanism to improve radiology report generation in terms of word-overlap metrics and time complexity. <br> • Utilized a pre-trained GPT2 conditioned on visual and weighted semantic features to promote faster training, eliminate vocabulary selection, and handle punctuation. <br> • The first study to employ semantic similarity metrics to quantitatively analyze medical report generation results. | • Conditioning mechanism tackled punctuations, vocabulary collection, and reduced training duration. <br> • The architecture does not require modification to be trained on distinct data sets. <br> • Incorporating semantic similarity in addition to word overlap metrics improved medical report evaluation. <br> $^4$ The model's generalization ability and robustness against over-fitting are both hindered when the size of the dataset is small. |
| CGI [367] | • Provides cohesive and precise X-ray reports in a fully differentiable manner by dividing the report generation system into a classifier, generator, and interpreter. <br> • Their conducted experiments revealed that incorporating additional scans besides clinical history can be beneficial in providing higher-quality X-ray reports. | • Flexibility in processing additional input data, such as clinical documents and extra scans, which also contributes to performance improvement. <br> • The model doesn't provide vital information, such as illness orientation and time-series correlations, which facilitates more reliable reports. |
| CGRG [368] | • Presented a Transformer-based method that estimates report uncertainty to develop a more reliable MRG system and facilitate diagnostic decision-making. <br> • Introduced the Sentence Matched Adjusted Semantic Similarity (SMAS) to capture vital and relevant features in radiology report generation more effectively. | • Assessing visual and textual uncertainties leads to more reliable reports in medical diagnosis. <br> • Measuring uncertainties can properly provide correlated confidence between various reports, which is beneficial to aiding radiologists in clinical report generation [392]. |
| CXR-RePaiR [369] | • Employs contrastive language image pre-training (CLIP) to generate chest X-ray reports. <br> • The utilization of representations pairs of X-ray reports aids the process of retrieving unstructured radiology reports written in natural language. | • Accuracy is evaluated by comparing it with the ground truth using a labeler, and performance scores are calculated accordingly. <br> • Their experiments indicate that despite generating precise diagnostic explanations, it may not consistently employ identical wording as the initial report. |
| BioViL-T [370] | • Provides a hybrid design that aligns X-ray report pairs using temporal information and prior knowledge. <br> • Combines CNN and transformer encoders to extract image representations and matches them with text representations using a chest X-ray domain-specific transformer model. <br> • Introduced MS-CXR-T, a benchmark dataset that facilitates the evaluation of vision-language representations in capturing temporal semantics. | • Prior report is more crucial for optimal performance compared to prior image as it outlines the image content and enhances the signal clarity. <br> • A significant factor in effective language supervision during the pre-training phase is the utilization of contrastive loss. |
| ChatCAD [371] | • Integrates MRG with large language models such as ChatGPT to improve diagnosis accuracy <br> • The LLM summarizes and corrects errors in the generated reports via the concatenation of textual representations from CAD networks. | • The experiments show that model size and complexity have a significant impact on generation accuracy, as larger models generally result in improved F1-scores. <br> • Due to the reports being less similar to human-written text, the system's BLEU score is lower compared to previous methods. |

## 9.3. Memory-based Systems

Concerning the development of systems that rely on a memory mechanism to generate medical reports, Chen et al. [357] presented a Memory-Driven Transformer (**MDT**), a model suitable for the generation of long informative reports and one of the first works on the MIMIC-CXR dataset [385]. MDT employs a relational memory to exploit characteristics prevailing in reports of similar images, and then the memory is incorporated into the decoder section of the Transformer by implementing a memory-driven conditional layer normalization (MCLN).

Likewise, Nooralahzadeh et al. [359] introduced **$M^2$ TR. progressive**, a report generation approach that utilizes curriculum learning, which is a strategy of training machine learning models by starting with easy samples and gradually increasing the samples' difficulty [395]. Instead of directly generating full reports from medical images, their work formulates the problem into two steps: first, the Meshed-Memory Transformer ($M^2$ TR.) [396], as a powerful image captioning model, receives the visual features extracted by a DenseNet [379] backbone and generates high-level global context. Second, BART [397], as a Transformer-based architecture, encodes these contexts with a bidirectional encoder and decodes its output using a left-to-right decoder into coherent reports. The overview of the process is depicted in Figure 42.
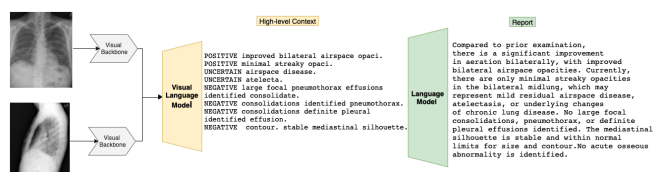


Figure 42: Workflow of the $M^2$ Tr. Progressive framework. The task is accomplished in two stages: First, the Meshed-Memory Transformer ($M^2$ TR.) receives visual features extracted by a DenseNet [379] backbone and generates high-level global context. Second, the BART [397] architecture encodes contexts with a bidirectional encoder and decodes them with a left-to-right decoder to produce coherent reports [359].

Additionally, You et al. [358] proposed **AlignTransformer**, a framework composed of two modules: Align Hierarchical Attention (AHA) and Multi-Grained Transformer (MGT). In their approach, first visual features and disease tags are extracted from the medical image by an image encoder, then they get aligned hierarchically to obtain multi-grained disease-grounded visual features in the AHA module. The obtained grounded features are capable of tackling the data bias problem by promoting a better representation of abnormal sections. Next, these grounded visual features are exploited by an adaptive exploiting attention (AEA) [396] mechanism in the MGT module for

Table 20: Comparison of Transformer-based medical report generation systems in terms of NLG performance metrics. The methods are ordered by BL-4 score, which captures more precise phrase and sentence structure.

| IU Chest X-ray [374] | | | | | | | |
|---|---|---|---|---|---|---|---|
| **Method** | **BL-1** | **BL-2** | **BL-3** | **BL-4** | **METEOR** | **ROUGE** | **CIDEr** |
| **RTMIC** [355] | 0.350 | 0.234 | 0.143 | 0.096 | - | - | 0.323 |
| **CDGPT2** [366] | 0.387 | 0.245 | 0.166 | 0.111 | 0.164 | 0.289 | 0.257 |
| **KERP** [362] | 0.482 | 0.325 | 0.226 | 0.162 | - | 0.339 | 0.280 |
| **MDT** [357] | 0.470 | 0.304 | 0.219 | 0.165 | 0.187 | 0.371 | - |
| **PPKED** [363] | 0.483 | 0.315 | 0.224 | 0.168 | 0.190 | 0.376 | **0.351** |
| **CMN** [361] | 0.475 | 0.309 | 0.222 | 0.170 | 0.191 | 0.375 | - |
| **AlignTransformer** [358] | 0.484 | 0.313 | 0.225 | 0.173 | 0.204 | 0.379 | - |
| **M$^2$ TR. progressive** [359] | 0.486 | 0.317 | 0.232 | 0.173 | 0.192 | 0.390 | - |
| **CGRG** [368] | **0.497** | **0.357** | **0.279** | **0.225** | **0.217** | **0.408** | - |
| MIMIC-CXR [385] | | | | | | | |
| **Method** | **BL-1** | **BL-2** | **BL-3** | **BL-4** | **METEOR** | **ROUGE** | **CIDEr** |
| **MDT** [357] | 0.353 | 0.218 | 0.145 | 0.103 | 0.142 | 0.277 | - |
| **PPKED** [363] | 0.360 | 0.224 | 0.149 | 0.106 | 0.149 | 0.284 | 0.237 |
| **CMN** [361] | 0.353 | 0.218 | 0.148 | 0.106 | 0.142 | 0.278 | - |
| **M$^2$ TR. progressive** [359] | 0.378 | 0.232 | 0.154 | 0.107 | 0.145 | 0.272 | - |
| **AlignTransformer** [358] | 0.378 | 0.235 | 0.156 | 0.112 | 0.158 | 0.283 | - |
| **CRG** [364] | 0.415 | 0.272 | 0.193 | 0.146 | 0.159 | 0.318 | - |
| **CGI** [367] | **0.495** | **0.360** | **0.278** | **0.224** | **0.222** | **0.390** | - |

the generation of the medical reports. They also justified their model's efficiency through the manual evaluation of clinical radiologists.

In **MDT-WCL** [360], the problem is approached with a weakly supervised contrastive loss, which lends more weight to the reports that are semantically close to the target reports, and a memory-driven Transformer is adopted as the backbone model to store key information in its memory module. To aid the contrastive learning during training, after clustering the reports into groups with the K-Means algorithm, each report is assigned a label corresponding to its cluster, and the semantically closed ones are considered to be in the same cluster.

Although previous approaches have achieved promising results, they lack the ability to generate mappings between images and texts to align visual-textual information and assist medical diagnosis. In order to facilitate visual-textual alignment, the Cross-modal Memory Network (**CMN**) [361] extended encoder-decoder methods by utilizing a shared memory for better alignment of information between images and texts. It uses a pre-trained ResNet [64] as the visual extractor to output visual features, then passes them to the cross-modal memory network that utilizes a matrix to store information where each row represents the embedding of information linking images and texts. To access the stored information aligning the modalities, memory querying and responding are implemented in a multi-threaded manner.

### 9.4. Other Systems

Other MRG systems focus on solving the problem with different ideas. Lovelace et al. [364] proposed a generation framework composed of two stages. In the first stage, a Transformer model is adopted to map the input image features extracted by a DenseNet-121 [379] to contextual annotations and learn report generation. In the second stage, a procedure is introduced to differentiably sample a clinical report from the Transformer decoder and obtain observational clinical information from the sample. This differentiability is further employed to fine-tune the model for improving clinical coherence by applying their differentiable CheXpert to the sampled reports. Fueled by recent progress in explainable artificial intelligence and the introduction of algorithms that attempt to provide interpretable prediction in DL-based systems, Likewise, in **CDGPT2** [366], the medical image is passed into a Chexnet [398] to provide localizations of 14 types of diseases from the images as visual features. To implement better semantic features, the model was fine-tuned as a multi-label classification problem to extract

manual tags from the IU-Xray dataset [374] by replacing the final layer of the model with a layer containing 105 neurons to produce 105 tags. The vector representation of the tags is then fed into a pre-trained distilGPT2 [399] as the decoder to generate medical reports. Moreover, Wang et al. [368] presented a confidence-guided report generation (**CGRG**) approach to support reliability in report generation by quantifying visual and textual uncertainties. It's comprised of an auto-encoder that reconstructs images, a Transformer encoder that encodes the input visual feature extracted by ResNet-101 [64], and a Transformer decoder for report generation. Visual uncertainty is obtained by the AutoEncoder, which acts as a guide for the visual feature extractor, and textual uncertainty is quantified based on the introduced Sentence Matched Adjusted Semantic Similarity (SMAS) which captures the similarity between the generated reports. These uncertainties are further utilized to aid the model optimization process.

The recent outbreak of COVID-19, one of the deadliest pandemics, has influenced the research community to alleviate the tedious and time-consuming work of producing medical reports. VL-BERT, [400] as an extension of BERT, [30] can be employed as an intelligent medical report generation system to expedite the diagnosis process. **Medical-VLBERT** [365] introduced VL-BERT to the medical report generation domain. It defines the problem as a two-step procedure: First, it utilizes two distinct VL-BERTs as terminology encoders to produce terminology-related features (textual and visual), and then these features are fed into a shared language decoder to produce medical textbooks and reports. The proposed method takes into account predefined terminology word embeddings that represent medical domain knowledge. These embeddings are paired distinctly with two other embeddings as an input to the encoders: textbook embeddings, which are generated by employing a lookup table, and spatial feature embeddings (termed "visual context") that are extracted from medical images by implementing DenseNet-121 [379]. The encoders then integrate this pairwise information separately to produce textual and visual terminological features. Subsequently, a shared language decoder is trained by utilizing an alternate approach to properly exchange the knowledge captured by the encoders.

Furthermore, in the work of Nguyen et al. [367], a classification, generation, and interpretation framework (**CGI**) is proposed to address clinical accuracy. Each term of the framework's name represents a different module to perform the task. The classification module learns how to discover diseases and generate their embeddings, which consist of an image and text encoder to extract the global visual features from medical images and obtain text-summarized embeddings from clinical documents. The generation module is a Transformer model that takes the disease embeddings as input and generates medical reports from them. The interpretation module then takes these reports for evaluation and fine-tuning.

Deep learning methods in medical report analysis require extensive amounts of annotated training data, which requires a significant amount of time and effort. To this end, self-supervised report generation approaches use a large amount of unlabeled text data without explicit supervision. For instance,

**CXR-RePaiR** [369] utilizes self-supervised contrastive language image pre-training. It encodes reports with the CheXbert transformer network[389], and leverages the X-ray report pair representations to retrieve unstructured radiology reports with a contrastive scoring function trained on chest X-ray-report pairs to rank the similarity between a test dataset of X-rays and a large corpus of reports. Likewise, **BioViL-T** [370] is a hybrid design that improves the alignment of X-ray report pairs by incorporating temporal information and prior knowledge to support complementary self-supervision and exploit the strengths of each modality. It compares prior images with given reports to exploit temporal correlations and enhance representation learning. BioViL-T uses a multi-image encoder to handle the absence of prior images and spatial misalignment, benefiting both image and text models. Image representations are extracted utilizing a hybrid CNN and transformer encoder, which are then matched with corresponding text representations obtained with CXR-BERT [401] by training with contrastive objectives.

Recently, large language models (LLMs), which are AI systems that learn the patterns of human language from massive amounts of text data, have revolutionized human-like text generation and natural language understanding. Consequently, they hold great potential in clinical applications, such as understanding medical texts and report generation [402, 403]. One of the successful large language models based on the GPT-3.5 architecture is ChatGPT [404], with extensive knowledge and abilities to perform a wide range of tasks such as question answering, text summarization, and generating textual instructions. Wang et al. [371] presented a scheme called **ChatCAD** that integrates LLMs such as GPT-3 [405] and ChatGPT with Computer-Aided Diagnosis (CAD) models to perform medical report generation. CAD systems are computer programs that assist doctors in interpreting medical images, such as X-rays or MRI scans. To generate reports via ChatCAD, the medical image is passed to a segmentation, image classification, and a report generation network. Since the outputs of segmentation and classification networks are a mask and a vector, respectively, they are transformed into textual format to be understandable by LLMs. These textual representations are then concatenated and presented to the LLM, which then summarizes the results from all the CAD networks and exploits them to correct the errors in the generated report. Their experiments indicate that LLMs can be more effective than conventional CAD systems in enhancing medical report quality.

> ### 9.5. Discussion and Conclusion
>
> This section offers a systematic review of the Transformer architectures configured for medical report generation. Compared to previous sections that reviewed ViT-based frameworks to tackle different medical tasks and problems, this section focuses mostly on using standard Transformers as the core of a medical report generation system. A common theme prevailing in these systems is to solve the problem with an encoder-decoder architecture supported by a CNN-based visual backbone. As mentioned in previous sections, the self-attention mechanism undermines

the representation of low-level details. On the other hand, since medical reports consist of long and multiple sentences, Transformers are of great significance to model long-term dependencies, which assists clinically accurate report generation [406, 367]. To exploit the power of both CNNs and Transformers simultaneously, state-of-the-art MRG systems usually embed CNNs along with Transformers in their frameworks [355, 366, 368]. We have provided information in Table 18 on the reviewed report generation methods concerning their architectural type, modality, organ, pre-trained strategy, datasets, metrics, and year. Table 19 contains summarized information about the methodologies, including their contributions and highlights. In addition, it should be noted that several survey publications have been published in this field of medicine [375, 390, 407], and the most recent one provided a technical overview of Transformer-based clinical report generation [22]. We approach our review differently by distinguishing the proposed methods based on the mechanism they used to support the prevailing concerns such as long and coherent text generation, reliability, and visual-textual biases.

The ultimate goal of these frameworks is to increase clinical accuracy to expedite the diagnosis process and reduce the workloads in radiology professions [364, 367]. Numerous works have attempted to facilitate diagnostic decision-making by aligning correlated sections of medical image and textual report that provide valuable information for detecting abnormalities [358, 361]. Also, multiple studies emphasized the importance of universal knowledge, and designed a system to incorporate prior information for detecting disease [362, 357]. Some research effort was also put into better representation learning by contrasting normal and abnormal samples against each other in representation space by utilizing a contrastive loss as the objective [360]. One recent work was inspired by curriculum learning to imitate the order of the human learning process [359].

As for the cumbersomeness of accessing annotated textual training data, research in MRG also focused on consolidating self-supervised learning and transformer architectures for report generation, which are discussed in Section 9.4 through the analysis of two specific works [408, 369]. Chat-CAD [371] was also introduced recently to integrate MRG with large language models like ChatGPT to improve diagnosis accuracy, but it achieves a lower BLEU score as it generates reports that are less similar to human-written text. The Table 20 includes performance values of BLEU, METEOR, ROUGE, and CIDEr for each method. Among the methods evaluated on the IU Chest X-ray, KERP and CGRG perform the best across most metrics, while RTMIC and CDGPT2 are the weakest. MDT has the highest METEOR score, CGRG achieves the highest ROUGE score, and PP-KED demonstrates the highest CIDEr score. In contrast, for the MIMIC-CXR dataset, CRG and CGI demonstrate the best performance, while MDT shows the weakest overall performance.

Additionally, Table 20 compares the performance of the re-

viewed methods Natural Language Generation (NLG) metrics. The methods are presented in separate sections based on the dataset they were evaluated on, with one section dedicated to the IU Chest X-ray [374], and the other to the MIMIC-CXR [385].

Overall, we believe that MRG systems need more research and progression to be robustly incorporated in a practical setting.

## 10. Open Challenges and Future Perspectives

So far, we discussed the application of Transformers (especially vision Transformers) and reviewed state-of-the-art models in medical image analysis. Even though their effectiveness is exemplified in previous sections by delicately presenting their ideas and analyzing the significant aspects that were addressed in their proposed methods, there is still room for improvement in many areas to devise a more practical and medically accurate system by leveraging Transformers. Consequently, we discuss the challenges and future directions hoping to help researchers gain insight into the limitations and develop more convenient automatic medical systems based on Transformers.

### 10.1. Explainability

Fueled by recent progress in XAI (explainable artificial intelligence) and the introduction of algorithms that attempt to provide interpretable prediction in DL-based systems, researchers are putting effort into incorporating XAI methods into constructing Transformer-based models to promote a more reliable and understandable system in different areas, including medical analysis [409, 410]. Existing approaches usually highlight important regions of the medical image that contribute to the model prediction by employing attention maps [411, 38]. Furthermore, Vision Transformers (ViTs) have the ability to provide attention maps that indicate the relevant correlations between the regions of the input and the prediction. However, the challenge of numerical instabilities in using propagation-based XAI methods such as LRP [412] and the vagueness of the attention maps, which leads to inaccurate token associations [76, 413], makes interpretable ViTs an open research opportunity in computer vision, especially in medical image analysis. We believe that including interpretable vision Transformers, such as ViT-NeT [413], in various medical applications can promote user-friendly predictions and facilitate decision-making in the diagnosis of medical conditions, and is a promising direction in medical research problems.

### 10.2. Richer Feature Representation

An effective and suitable representation space is substantially influential in building medical analysis systems. Transformers have demonstrated their efficiency in obtaining global information and capturing long-term dependencies in many areas, such as Natural Language Processing (NLP), Computer Vision, and Speech Recognition [335], and CNNs have proven to be effective in extracting local context from visual data [414]. However,

this locality usually enables these networks to capture rich local texture representation and lacks model global dependency. As a result, many approaches stack Transformers along with CNNs to leverage both local and global information simultaneously in clinical applications (e.g., medical report generation) [358, 364, 48]. Recent studies stated that the single-scale representation of ViTs hinders improvement in dense prediction tasks, so a multi-scaled feature representation is implemented which achieves better performance in computer vision tasks, including image classification, object detection, and image segmentation [415, 416]. Generalizing this idea to medical applications of ViTs to facilitate devising a clinically suitable system can be considered as future work.

### 10.3. Video-based analysis

There has been an increasing interest in the vision community in extending ViT architectures to video recognition tasks. Recently, a handful of papers have integrated standard Transformers with their models in AI-assisted dynamic clinical tasks [417, 418, 419, 420]. However, the scarcity of the proposed approaches puts video-based medical analysis in an infancy stage and open for future investigations. Another potential research direction is to explore the power of video vision Transformer variants, such as Video Swin Transformer [421], in clinical video understanding and to facilitate automatic robotic surgery.

### 10.4. High Computational Complexity

The robustness of Transformer models in layouts that implement large numbers of parameters is one of their strengths. While this is a beneficial trait that makes it possible to train models of enormous scale, it leads to the requirement of large resources for training and inferencing [21]. Particularly disadvantageous to medical image analysis is that expanding the use of ViTs for pretraining in new tasks and datasets comes with substantial expenses and burdens. Additionally, gathering medical samples can be difficult and the dataset scale is often limited. For instance, according to empirical studies in [15], pretraining a ViT-L/16 model on the large-scale dataset of ImageNet takes approximately 30 days employing a standard cloud TPUv3 with 8 cores. As a result, a notable number of papers utilized the pre-trained weights of ViT models to exploit the transfer learning strategy to alleviate training load [41, 17, 40], but in some cases, such as dealing with volumetric medical images, where transfer learning doesn't demonstrate any improvements [159, 338], the pretraining process is necessary to capture domain-specific features for generalization and better performance.

In addition, the standard self-attention mechanism suffers from a significant drawback: its computational complexity grows quadratically with the number of tokens ($n$), making it impractical for tasks where input images can be millions of pixels in size or when dealing with volumetric data. To address this issue, several approaches have been proposed to reduce the computational burden of self-attention. They either reduce the number of tokens by windowing [422, 58], shift the calculation to the channel dimension [423, 424], or change the order of multiplying query, key, and value [425]. In the context

of 3D biomedical image segmentation, Zhang et al. [426] introduced a dynamic linear Transformer algorithm that achieves linear complexity by focusing computations only on the region of interest (ROI). This significantly reduces the overall computational requirements. Shen et al. [425] proposed an efficient attention mechanism that avoids the pairwise similarity computations of dot-product attention. Instead, it normalizes the keys and queries first, performs multiplication between the keys and values, and then multiplies the resulting global context vectors with the queries. This approach reduces the computational complexity of self-attention from $O(dn^2)$ to $O(d^2n)$, where $d$ is the embedding dimension. XCiT, [423] another approach, addresses the complexity challenge by introducing cross-covariance attention, which has a complexity of $O(\frac{d^2n}{h})$, where $h$ is the number of attention heads. EdgeNext, an architecture proposed by Maaz et al. [424] for edge devices, optimizes the number of Multiplication-Addition (MAdd) operations required. It employs a variation of the attention mechanism used in XCiT called split depth-wise transpose attention (SDTA). Huang et al. [154] introduce efficient self-attention (ESA), specifically for the spatial reduction in 3D data. ESA reduces the number of tokens by a spatial reduction ratio $R$ while expanding the channel dimension by the same ratio. As a result, the complexity of ESA is reduced to $O(\frac{n^2}{R})$.

Despite progress in making self-attention more efficient for computer vision tasks, further work is needed to develop efficient Transformer models tailored to the medical domain, particularly for handling challenging 3D data or large-scale images (e.g., WSI). Moreover, the ultimate goal is to deploy these models on low-cost edge devices or in resource-constrained environments. Therefore, the focus should be on designing Transformer systems that strike a balance between computational complexity and clinical accuracy while ensuring robustness. This area of research holds great promise and should be pursued further.

### 10.5. Transformer-based Registration

As reviewed in Section 8, the idea of employing Transformers to support efficient medical image registration has become popular in recent years. The ability of the self-attention mechanism assists the learning of long-term visual correlations since their unlimited receptive field promotes a more accurate understanding of the spatial relationship between moving and fixed images [338, 339]. However, registration systems composed of Transformer architectures are still in their infancy and require more research effort to be put into them.

### 10.6. Data-Driven Predictions

With supervised learning as a popular fashion in building intelligent systems, the model learns features based on the provided annotations that are suitable to accomplish a specific task, which hinders generalizability. In other words, supervised learning modifies the bias-variance trade-off in favor of the strong inductive biases that lead to making assumptions as a means to aid the model in learning a particular task quicker and with higher sample efficiency. However, these hard assumptions sacrifice adaptability to other settings and unseen datasets,

and the model learns to accomplish its task without having an innate understanding of the data. To tackle this issue, unsupervised regimes enable the algorithms to act as general descriptors and capture features that will assist them in performing efficiently in a wide range of tasks. Similarly, in medical image analysis, adopting Transformer networks with unsupervised learning algorithms promotes robustness and generalizability to other datasets and tasks.

### 10.7. Medical Software Ecosystems

A future direction for advancing in the automatic medical analysis is to provide an open-source environment that contains libraries suitable for solving multiple medical tasks and challenges with Transformer architectures. Developers can further contribute to the ecosystem by updating and adding additional tasks, bringing novelty, and proposing ideas to enhance performance and accuracy [152]. Companies and organizations can support the system by preparing the necessary computational resources and hardware requirements. Sample of software prototypes in this direction are nnU-Net [427], Ivadomed [428], and preliminary works such as [166], which provides an end-to-end pipeline for implementing deep models on medical data.

## 11. Discussion and Conclusion

In this paper, we presented a comprehensive encyclopedic review of the applications of Transformers in medical imaging. First, we provided preliminary information regarding the Transformer structures and the idea behind the self-attention mechanism in the introduction and background sections. Starting from Section 3, we reviewed the literature on Transformer architecture in diverse medical imaging tasks, namely, classification, segmentation, detection, reconstruction, synthesis, registration, and clinical report generation. For each application, we provided a taxonomy and high-level abstraction of the core techniques employed in these models along with the SOTA approaches. We also provided comparison tables to highlight the pros and cons, network parameters, type of imaging modality they are considering, organ, and the metrics they are using. Finally, we outlined possible avenues for future research directions.

## References

[1] J. Arevalo, F. A. González, R. Ramos-Pollán, J. L. Oliveira, M. A. G. Lopez, Representation learning for mammography mass lesion classification with convolutional neural networks, Computer methods and programs in biomedicine 127 (2016) 248–257.

[2] S. Karimijafarbigloo, R. Azad, A. Kazerouni, D. Merhof, MS-Former: Multi-scale self-guided transformer for medical image segmentation, in: Medical Imaging with Deep Learning, 2023.

[3] R. Azad, M. Asadi-Aghbolaghi, M. Fathy, S. Escalera, Bi-directional ConvLSTM U-Net with densley connected convolutions, in: 2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW), 2019, pp. 406–415.

[4] R. Azad, N. Khosravi, D. Merhof, SMU-Net: Style matching U-Net for brain tumor segmentation with missing modalities, in: International Conference on Medical Imaging with Deep Learning, PMLR, 2022, pp. 48–62.

[5] P. Ramachandran, N. Parmar, A. Vaswani, I. Bello, A. Levskaya, J. Shlens, Stand-alone self-attention in vision models, Advances in Neural Information Processing Systems 32 (2019).

[6] I. Bello, B. Zoph, A. Vaswani, J. Shlens, Q. V. Le, Attention augmented convolutional networks, in: Proceedings of the IEEE/CVF international conference on computer vision, 2019, pp. 3286–3295.

[7] A. Vaswani, P. Ramachandran, A. Srinivas, N. Parmar, B. Hechtman, J. Shlens, Scaling local self-attention for parameter efficient visual backbones, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 12894–12904.

[8] X. Wang, R. Girshick, A. Gupta, K. He, Non-local neural networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 7794–7803.

[9] J. Hu, L. Shen, G. Sun, Squeeze-and-excitation networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 7132–7141.

[10] M. Al-Shabi, K. Shak, M. Tan, ProCAN: Progressive growing channel attentive non-local network for lung nodule classification, Pattern Recognition 122 (2022) 108309.

[11] D. V. Sang, T. Q. Chung, P. N. Lan, D. V. Hang, D. Van Long, N. T. Thuy, Ag-curesnest: A novel method for colon polyp segmentation, arXiv preprint arXiv:2105.00402 (2021).

[12] C. Yao, J. Tang, M. Hu, Y. Wu, W. Guo, Q. Li, X.-P. Zhang, Claw U-Net: A UNet variant network with deep feature concatenation for scleral blood vessel segmentation, in: CAAI International Conference on Artificial Intelligence, Springer, 2021, pp. 67–78.

[13] T. Gonçalves, I. Rio-Torto, L. F. Teixeira, J. S. Cardoso, A survey on attention mechanisms for medical applications: are we moving towards better algorithms? (2022).

[14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, Advances in neural information processing systems 30 (2017).

[15] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An image is worth 16x16 words: Transformers for image recognition at scale, arXiv preprint arXiv:2010.11929 (2020).

[16] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, J. Dai, Deformable {detr}: Deformable transformers for end-to-end object detection, in: International Conference on Learning Representations, 2021.
URL https://openreview.net/forum?id=gZ9hCDWe6ke

[17] J. Chen, Y. Lu, Q. Yu, X. Luo, E. Adeli, Y. Wang, L. Lu, A. L. Yuille, Y. Zhou, Transunet: Transformers make strong encoders for medical image segmentation, arXiv preprint arXiv:2102.04306 (2021).

[18] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, C. Schmid, Vivit: A video vision transformer, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 6836–6846.

[19] X. Chen, X. Wang, J. Zhou, Y. Qiao, C. Dong, Activating more pixels in image super-resolution transformer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 22367–22377.

[20] B. Azad, R. Azad, S. Eskandari, A. Bozorgpour, A. Kazerouni, I. Rekik, D. Merhof, Foundational models in medical imaging: A comprehensive survey and future vision, arXiv preprint arXiv:2310.18689 (2023).

[21] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, M. Shah, Transformers in vision: A survey, ACM computing surveys (CSUR) 54 (10s) (2022) 1–41.

[22] F. Shamshad, S. Khan, S. W. Zamir, M. H. Khan, M. Hayat, F. S. Khan, H. Fu, Transformers in medical imaging: A survey, Medical Image Analysis (2023) 102802.

[23] K. He, C. Gan, Z. Li, I. Rekik, Z. Yin, W. Ji, Y. Gao, Q. Wang, J. Zhang, D. Shen, Transformers in medical image analysis: A review, Intelligent Medicine (2022).

[24] K. S. Kalyan, A. Rajasekharan, S. Sangeetha, Ammus: A survey of transformer-based pretrained models in natural language processing, arXiv preprint arXiv:2108.05542 (2021).

[25] A. Kolesnikov, L. Beyer, X. Zhai, J. Puigcerver, J. Yung, S. Gelly, N. Houlsby, Big transfer (bit): General visual representation learning, in: European conference on computer vision, Springer, 2020, pp. 491–507.

[26] H. Cao, Y. Wang, J. Chen, D. Jiang, X. Zhang, Q. Tian, M. Wang, Swin-unet: Unet-like pure transformer for medical image segmentation, in: Proceedings of the European Conference on Computer Vision Workshops(ECCVW), 2022.

[27] R. Azad, M. Heidari, M. Shariatnia, E. K. Aghdam, S. Karimijafarbigloo, E. Adeli, D. Merhof, TransDeepLab: Convolution-free transformer-based deeplab v3+ for medical image segmentation, arXiv preprint arXiv:2208.00713 (2022).

[28] O. Ronneberger, P. Fischer, T. Brox, U-net: Convolutional networks for biomedical image segmentation, in: International Conference on Medical image computing and computer-assisted intervention, Springer, 2015, pp. 234–241.

[29] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in: Proceedings of the European conference on computer vision (ECCV), 2018, pp. 801–818.

[30] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, arXiv preprint arXiv:1810.04805 (2018).

[31] M. Heidari, A. Kazerouni, M. Soltany, R. Azad, E. K. Aghdam, J. Cohen-Adad, D. Merhof, Hiformer: Hierarchical multi-scale representations using transformers for medical image segmentation, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 2023, pp. 6202–6212.

[32] P. Benz, S. Ham, C. Zhang, A. Karjauv, I. S. Kweon, Adversarial robustness comparison of vision transformer and mlp-mixer to cnns, arXiv preprint arXiv:2110.02797 (2021).

[33] S. Park, G. Kim, J. Kim, B. Kim, J. C. Ye, Federated split task-agnostic vision transformer for COVID-19 CXR diagnosis, in: A. Beygelzimer, Y. Dauphin, P. Liang, J. W. Vaughan (Eds.), Advances in Neural Information Processing Systems, 2021.
URL https://openreview.net/forum?id=Ggikq6Tdxch

[34] K. Singhal, T. Tu, J. Gottweis, R. Sayres, E. Wulczyn, L. Hou, K. Clark, S. Pfohl, H. Cole-Lewis, D. Neal, et al., Towards expert-level medical question answering with large language models, arXiv preprint arXiv:2305.09617 (2023).

[35] L. Seenivasan, M. Islam, G. Kannan, H. Ren, SurgicalGPT: End-to-end language-vision GPT for visual question answering in surgery, arXiv preprint arXiv:2304.09974 (2023).

[36] X. Gao, Y. Qian, A. Gao, COVID-VIT: Classification of COVID-19 from CT chest images based on vision transformer models, arXiv preprint arXiv:2107.01682 (2021).

[37] S. Yu, K. Ma, Q. Bi, C. Bian, M. Ning, N. He, Y. Li, H. Liu, Y. Zheng, Mil-vt: Multiple instance learning enhanced vision transformer for fundus image classification, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 45–54.

[38] A. K. Mondal, A. Bhattacharjee, P. Singla, A. Prathosh, xViTCOS: explainable vision transformer based COVID-19 screening using radiography, IEEE Journal of Translational Engineering in Health and Medicine 10 (2021) 1–10.

[39] C. Matsoukas, J. F. Haslum, M. Söderberg, K. Smith, Is it time to replace cnns with transformers for medical images?, arXiv preprint arXiv:2108.09038 (2021).

[40] B. Gheflati, H. Rivaz, Vision transformers for classification of breast ultrasound images, in: 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), IEEE, 2022, pp. 480–483.

[41] D. Shome, T. Kar, S. N. Mohanty, P. Tiwari, K. Muhammad, A. AlTameem, Y. Zhang, A. K. J. Saudagar, Covid-transformer: Interpretable COVID-19 detection using vision transformer for healthcare, International Journal of Environmental Research and Public Health 18 (21) (2021) 11086.

[42] R. Yan, L. Qu, Q. Wei, S.-C. Huang, L. Shen, D. Rubin, L. Xing, Y. Zhou, Label-efficient self-supervised federated learning for tackling data heterogeneity in medical imaging, IEEE Transactions on Medical Imaging (2023).

[43] S. Perera, S. Adhikari, A. Yilmaz, POCFormer: A lightweight transformer architecture for detection of COVID-19 using point of care ultrasound, in: 2021 IEEE International Conference on Image Processing (ICIP), IEEE, 2021, pp. 195–199.

[44] M. Bhattacharya, S. Jain, P. Prasanna, Radiotransformer: a cascaded global-focal transformer for visual attention–guided disease classification, in: European Conference on Computer Vision, Springer, 2022, pp. 679–698.

[45] C. Liu, Q. Yin, Automatic diagnosis of COVID-19 using a tailored transformer-like network, in: Journal of Physics: Conference Series, Vol. 2010, IOP Publishing, 2021, p. 012175.

[46] Y. Dai, Y. Gao, F. Liu, Transmed: Transformers advance multi-modal medical image classification, Diagnostics 11 (8) (2021) 1384.

[47] S. Wang, Z. Zhuang, K. Xuan, D. Qian, Z. Xue, J. Xu, Y. Liu, Y. Chai, L. Zhang, Q. Wang, et al., 3DMET: 3D medical image transformer for knee cartilage defect assessment, in: International Workshop on Machine Learning in Medical Imaging, Springer, 2021, pp. 347–355.

[48] L. Tanzi, A. Audisio, G. Cirrincione, A. Aprato, E. Vezzetti, Vision transformer for femur fracture classification, Injury (2022).

[49] S. Park, G. Kim, Y. Oh, J. B. Seo, S. M. Lee, J. H. Kim, S. Moon, J.-K. Lim, J. C. Ye, Vision transformer for COVID-19 CXR diagnosis using chest X-ray feature corpus, arXiv preprint arXiv:2103.07055 (2021).

[50] R. Sun, Y. Li, T. Zhang, Z. Mao, F. Wu, Y. Zhang, Lesion-aware transformers for diabetic retinopathy grading, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 10938–10947.

[51] H. Li, F. Yang, Y. Zhao, X. Xing, J. Zhang, M. Gao, J. Huang, L. Wang, J. Yao, DT-MIL: Deformable transformer for multi-instance learning on histopathological image, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 206–216.

[52] Z. Shao, H. Bian, Y. Chen, Y. Wang, J. Zhang, X. Ji, et al., Transmil: Transformer based correlated multiple instance learning for whole slide image classification, Advances in Neural Information Processing Systems 34 (2021).

[53] S. Mehta, X. Lu, W. Wu, D. Weaver, H. Hajishirzi, J. G. Elmore, L. G. Shapiro, End-to-end diagnosis of breast biopsy images with transformers, Medical Image Analysis 79 (2022) 102466.

[54] Y. Zheng, R. H. Gindra, E. J. Green, E. J. Burks, M. Betke, J. E. Beane, V. B. Kolachalama, A graph-transformer for whole slide image classification, IEEE transactions on medical imaging 41 (11) (2022) 3003–3015.

[55] X. Huo, G. Sun, S. Tian, Y. Wang, L. Yu, J. Long, W. Zhang, A. Li, HiFuse: Hierarchical multi-scale feature fusion network for medical image classification, arXiv preprint arXiv:2209.10218 (2022).

[56] O. N. Manzari, H. Ahmadabadi, H. Kashiani, S. B. Shokouhi, A. Ayatollahi, MedViT: A robust vision transformer for generalized medical image classification, Computers in Biology and Medicine 157 (2023) 106791.

[57] A. Aminimehr, A. Molaei, E. Cambria, Entri: Ensemble learning with tri-level representations for explainable scene recognition, arXiv preprint arXiv:2307.12442 (2023).

[58] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, B. Guo, Swin transformer: Hierarchical vision transformer using shifted windows, in: Proceedings of the

IEEE/CVF International Conference on Computer Vision, 2021, pp. 10012–10022.

[59] Z. Xia, X. Pan, S. Song, L. E. Li, G. Huang, Vision transformer with deformable attention, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 4794–4803.

[60] M. Fayyaz, S. Abbasi Kouhpayegani, F. Rezaei Jafari, E. Sommerlade, H. R. Vaezi Joze, H. Pirsiavash, J. Gall, Adaptive token sampling for efficient vision transformers, European Conference on Computer Vision (ECCV) (2022).

[61] X. Dong, J. Bao, D. Chen, W. Zhang, N. Yu, L. Yuan, D. Chen, B. Guo, Cswin transformer: A general vision transformer backbone with cross-shaped windows, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 12124–12134.

[62] W. Li, X. Wang, X. Xia, J. Wu, X. Xiao, M. Zheng, S. Wen, Sepvit: Separable vision transformer, arXiv preprint arXiv:2203.15380 (2022).

[63] T. Yao, Y. Li, Y. Pan, Y. Wang, X.-P. Zhang, T. Mei, Dual vision transformer, IEEE transactions on pattern analysis and machine intelligence (2023).

[64] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.

[65] H. Touvron, M. Cord, M. Douze, F. Massa, A. Sablay-rolles, H. Jégou, Training data-efficient image transformers & distillation through attention, in: International Conference on Machine Learning, PMLR, 2021, pp. 10347–10357.

[66] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, A. Joulin, Emerging properties in self-supervised vision transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 9650–9660.

[67] M. H. Yap, G. Pons, J. Martí, S. Ganau, M. Sentis, R. Zwiggelaar, A. K. Davison, R. Marti, Automated breast ultrasound lesions detection using convolutional neural networks, IEEE journal of biomedical and health informatics 22 (4) (2017) 1218–1226.

[68] W. Al-Dhabyani, M. Gomaa, H. Khaled, A. Fahmy, Dataset of breast ultrasound images, Data in brief 28 (2020) 104863.

[69] X. Qi, L. G. Brown, D. J. Foran, J. Nosher, I. Haci-haliloglu, 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) (2021) 197–206.

[70] W. El-Shafai, F. Abd El-Samie, Extensive COVID-19 X-ray and CT chest images dataset, Mendeley data 3 (10) (2020).

[71] U. Sait, K. Lal, S. Prajapati, R. Bhaumik, T. Kumar, S. Sanjana, K. Bhalla, Curated dataset for COVID-19 posterior-anterior chest radiography images (X-Rays), Mendeley Data 1 (2020).

[72] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-cam: Visual explanations from deep networks via gradient-based localization, in: Proceedings of the IEEE international conference on computer vision, 2017, pp. 618–626.

[73] J. Deng, A large-scale hierarchical image database, Proc. of IEEE Computer Vision and Pattern Recognition, 2009 (2009).

[74] H. Gunraj, COVIDx CT-2A: A large-scale chest CT dataset for COVID-19 detection (2021).

[75] J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. Ball, K. Shpan-skaya, et al., Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison, in: Proceedings of the AAAI conference on artificial intelligence, Vol. 33, 2019, pp. 590–597.

[76] H. Chefer, S. Gur, L. Wolf, Transformer interpretability beyond attention visualization, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 782–791.

[77] Asia Pacific Tele-Ophthalmology Society, Aptos 2019 blindness detection: Detect diabetic retinopathy to stop blindness before it's too late., https://www.kaggle.com/c/aptos2019-blindness-detection/ (2019).

[78] S. Pachade, P. Porwal, D. Thulkar, M. Kokare, G. Deshmukh, V. Sahasrabuddhe, L. Giancardo, G. Quellec, F. Mériaudeau, Retinal fundus multi-disease image dataset (RFMiD): a dataset for multi-disease detection research, Data 6 (2) (2021) 14.

[79] D. Kollias, A. Arsenos, L. Soukissian, S. Kollias, MIA-COV19D: COVID-19 detection through 3-D chest CT image analysis, in: 2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), IEEE, 2021, pp. 537–544.

[80] F. Iandola, M. Moskewicz, S. Karayev, R. Girshick, T. Darrell, K. Keutzer, Densenet: Implementing efficient convnet descriptor pyramids, arXiv preprint arXiv:1404.1869 (2014).

[81] S. Wang, B. Z. Li, M. Khabsa, H. Fang, H. Ma, Linformer: Self-attention with linear complexity, arXiv preprint arXiv:2006.04768 (2020).

[82] J. Born, G. Brändle, M. Cossio, M. Disdier, J. Goulet, J. Roulin, N. Wiedemann, POCOVID-Net: automatic detection of COVID-19 from a new lung ultrasound imaging dataset (POCUS), arXiv preprint arXiv:2004.12084 (2020).

[83] L. Yuan, Q. Hou, Z. Jiang, J. Feng, S. Yan, Volo: Vision outlooker for visual recognition, IEEE transactions on pattern analysis and machine intelligence 45 (5) (2022) 6575–6586.

[84] M. E. Chowdhury, T. Rahman, A. Khandakar, R. Mazhar, M. A. Kadir, Z. B. Mahbub, K. R. Islam, M. S. Khan, A. Iqbal, N. Al Emadi, et al., Can AI help in screening viral and COVID-19 pneumonia?, IEEE Access 8 (2020) 132665–132676.

[85] J. P. Cohen, P. Morrison, L. Dao, K. Roth, T. Duong, M. Ghassem, COVID-19 image data collection: Prospective predictions are the future, Machine Learning for Biomedical Imaging 1 (2020) 1–38. doi:https://doi.org/10.59275/j.melba.2020-48g7. URL https://melba-journal.org/2020:002

[86] S. Mehta, X. Lu, D. Weaver, J. G. Elmore, H. Hajishirzi, L. Shapiro, Hatnet: an end-to-end holistic attention network for diagnosis of breast biopsy images, arXiv preprint arXiv:2007.13007 (2020).

[87] A. Srinivas, T.-Y. Lin, N. Parmar, J. Shlens, P. Abbeel, A. Vaswani, Bottleneck transformers for visual recognition, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2021, pp. 16519–16529.

[88] J. Redmon, A. Farhadi, Yolov3: An incremental improvement, arXiv preprint arXiv:1804.02767 (2018).

[89] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 2818–2826.

[90] L. Tanzi, E. Vezzetti, R. Moreno, A. Aprato, A. Audisio, A. Massè, Hierarchical fracture classification of proximal femur X-Ray images using a multistage deep learning approach, European journal of radiology 133 (2020) 109373.

[91] M. Tan, Q. Le, Efficientnet: Rethinking model scaling for convolutional neural networks, in: International conference on machine learning, PMLR, 2019, pp. 6105–6114.

[92] World-Health-Organization, Breast cancer, https://www.who.int/news-room/fact-sheets/detail/breast-cancer (2021).

[93] M. Y. Lu, D. F. Williamson, T. Y. Chen, R. J. Chen, M. Barbieri, F. Mahmood, Data-efficient and weakly supervised computational pathology on whole-slide images, Nature biomedical engineering 5 (6) (2021) 555–570.

[94] Y. Sharma, A. Shrivastava, L. Ehsan, C. A. Moskaluk, S. Syed, D. Brown, Cluster-to-conquer: A framework for end-to-end multi-instance learning for whole slide image classification, in: Medical Imaging with Deep Learning, PMLR, 2021, pp. 682–698.

[95] N. Naik, A. Madani, A. Esteva, N. S. Keskar, M. F. Press, D. Ruderman, D. B. Agus, R. Socher, Deep learning-enabled breast cancer hormonal receptor status determination from base-level h&e stains, Nature communications 11 (1) (2020) 1–8.

[96] M. Ilse, J. Tomczak, M. Welling, Attention-based deep multiple instance learning, in: International conference on machine learning, PMLR, 2018, pp. 2127–2136.

[97] B. Li, Y. Li, K. W. Eliceiri, Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 14318–14328.

[98] G. Campanella, M. G. Hanna, L. Geneslaw, A. Miraflor, V. Werneck Krauss Silva, K. J. Busam, E. Brogi, V. E. Reuter, D. S. Klimstra, T. J. Fuchs, Clinical-grade computational pathology using weakly supervised deep learning on whole slide images, Nature medicine 25 (8) (2019) 1301–1309.

[99] W. Li, V.-D. Nguyen, H. Liao, M. Wilder, K. Cheng, J. Luo, Patch transformer for multi-tagging whole slide histopathology images, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2019, pp. 532–540.

[100] M. Combalia, N. C. Codella, V. Rotemberg, B. Helba, V. Vilaplana, O. Reiter, C. Carrera, A. Barreiro, A. C. Halpern, S. Puig, et al., Bcn20000: Dermoscopic lesions in the wild, arXiv preprint arXiv:1908.02288 (2019).

[101] R. S. Lee, F. Gimenez, A. Hoogi, K. K. Miyake, M. Gorovoy, D. L. Rubin, A curated mammography data set for use in computer-aided detection and diagnosis research, Scientific data 4 (1) (2017) 1–9.

[102] M. D. L. I. Vayá, J. M. Saborit, J. A. Montell, A. Pertusa, A. Bustos, M. Cazorla, J. Galant, X. Barber, D. Orozco-Beltrán, F. García-García, et al., BIMCV COVID-19+: a large annotated dataset of RX and CT images from COVID-19 patients, arXiv preprint arXiv:2006.01174 (2020).

[103] A. Signoroni, M. Savardi, S. Benini, N. Adami, R. Leonardi, P. Gibellini, F. Vaccher, M. Ravanelli, A. Borghesi, R. Maroldi, et al., BS-Net: Learning COVID-19 pneumonia severity on a large chest X-ray dataset, Medical Image Analysis 71 (2021) 102046.

[104] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, R. M. Summers, Chestx-ray8: Hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 2097–2106.

[105] SIIM-ACR, SIIM-ACR Pneumothorax Segmentation, https://www.kaggle.com/c/siim-acr-pneumothorax-segmentation (2019).

[106] RSNA, RSNA Pneumonia Detection Challenge, https://www.kaggle.com/c/rsna-pneumonia-detection-challenge (2018).

[107] R. S. of North America, RSNA Pneumonia Detection Challenge, https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/ (2018).

[108] D. S. Kermany, M. Goldbaum, W. Cai, C. C. Valentim, H. Liang, S. L. Baxter, A. McKeown, G. Yang, X. Wu, F. Yan, et al., Identifying medical diagnoses and treatable diseases by image-based deep learning, Cell 172 (5) (2018) 1122–1131.
URL https://www.kaggle.com/datasets/paultimothymooney/chest-xray-pneumonia

[109] T. Rahman, A. Khandakar, Y. Qiblawey, A. Tahir, S. Kiranyaz, S. B. A. Kashem, M. T. Islam, S. Al Maadeed, S. M. Zughaier, M. S. Khan, et al., Exploring the effect of image enhancement techniques on COVID-19 detection using chest X-ray images, Computers in biology and medicine 132 (2021) 104319.

[110] H. Q. Nguyen, K. Lam, L. T. Le, H. H. Pham, D. Q. Tran, D. B. Nguyen, D. D. Le, C. M. Pham, H. T. Tong, D. H. Dinh, et al., Vindr-cxr: An open dataset of chest X-rays with radiologist's annotations, Scientific Data 9 (1) (2022) 1–7.

[111] P. Lakhani, J. Mongan, C. Singhal, Q. Zhou, K. P. Andriole, W. F. Auffermann, P. Prasanna, T. Pham, M. Peterson, P. J. Bergquist, et al., The 2021 SIIM-FISABIO-RSNA machine learning COVID-19 challenge: Annotation and standard exam classification of COVID-19 chest radiographs. (2021).

[112] E. B. Tsai, S. Simpson, M. P. Lungren, M. Hershman, L. Roshkovan, E. Colak, B. J. Erickson, G. Shih, A. Stein, J. Kalpathy-Cramer, et al., The RSNA international COVID-19 open radiology database (RICORD), Radiology 299 (1) (2021) E204–E213.

[113] E. B. Tsai, S. Simpson, M. P. Lungren, M. Hershman, L. Roshkovan, E. Colak, B. J. Erickson, G. Shih, A. Stein, J. Kalpathy-Cramer, et al., Data from medical imaging data resource center (MIDRC) - RSNA international COVID radiology database (RICORD) release 1c - chest X-ray, covid+ (MIDRC-RICORD-1c), The Cancer Imaging Archive (2021).

[114] K. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, S. Moore, S. Phillips, D. Maffitt, M. Pringle, et al., The cancer imaging archive (TCIA): maintaining and operating a public information repository, Journal of digital imaging 26 (6) (2013) 1045–1057.

[115] J. Saltz, M. Saltz, P. Prasanna, R. Moffitt, J. Hajagos, E. Bremer, J. Balsamo, T. Kurc, Stony brook university COVID-19 positive cases [data set] (2021). doi:10.7937/TCIA.BBAG-2923.

[116] EyePACKS, Kaggle diabetic retinopathy detection competition., https://www.kaggle.com/c/diabetic-retinopathy-detection (2015).

[117] N. C. Codella, D. Gutman, M. E. Celebi, B. Helba, M. A. Marchetti, S. W. Dusza, A. Kalloo, K. Liopyris, N. Mishra, H. Kittler, et al., Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic), in: 2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018), IEEE, 2018, pp. 168–172.

[118] N. Codella, V. Rotemberg, P. Tschandl, M. E. Celebi, S. Dusza, D. Gutman, B. Helba, A. Kalloo, K. Liopyris, M. 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).

[119] V. Rotemberg, N. Kurtansky, B. Betz-Stablein, L. Caffery, E. Chousakos, N. Codella, M. Combalia, S. Dusza, P. Guitera, D. Gutman, et al., A patient-centric dataset of images and metadata for identifying melanomas using clinical context, Scientific data 8 (1) (2021) 34.

[120] B. E. Bejnordi, M. Veta, P. J. Van Diest, B. Van Ginneken, N. Karssemeijer, G. Litjens, J. A. Van Der Laak, M. Hermsen, Q. F. Manson, M. Balkenhol, et al., Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer, Jama 318 (22) (2017) 2199–2210.

[121] B. Albertina, M. Watson, C. Holback, R. Jarosz, S. Kirk, Y. Lee, K. Rieger-Christ, J. Lemmerman, The cancer genome atlas lung adenocarcinoma collection (TCGA-LUAD) (version 4) [data set] (2016). doi:10.7937/K9/TCIA.2016.JGNIHEP5.

[122] S. Kirk, Y. Lee, P. Kumar, J. Filippini, B. Albertina, M. Watson, K. Rieger-Christ, J. Lemmerman, The cancer genome atlas lung squamous cell carcinoma collection (TCGA-LUSC) (version 4) (2016). doi:10.7937/K9/TCIA.2016.TYGKKFMQ.

[123] S. AB, V. R, J. C, A. O, K. J, H. E, F. J, S. NI, S. CA, B. TK, R. DL, O. A, H. MT, S. VR, K. V, S. SG, Radiogenomics of clear cell renal cell carcinoma: Preliminary findings of the cancer genome atlas-renal cell carcinoma (tcga-rcc) research group (2014). doi:10.7937/K9/TCIA.2014.K6M61GDW.

[124] E. Decencière, X. Zhang, G. Cazuguel, B. Lay, B. Cochener, C. Trone, P. Gain, R. Ordonez, P. Massin, A. Erginay, et al., Feedback on a publicly distributed image database: the messidor database, Image Analysis & Stereology 33 (3) (2014) 231–234.

[125] J. Krause, V. Gulshan, E. Rahimy, P. Karth, K. Widner, G. S. Corrado, L. Peng, D. R. Webster, Grader variability and the importance of reference standards for evaluating machine learning models for diabetic retinopathy, Ophthalmology 125 (8) (2018) 1264–1272.

[126] N. Bien, P. Rajpurkar, R. L. Ball, J. Irvin, A. Park, E. Jones, M. Bereket, B. N. Patel, K. W. Yeom, K. Shpanskaya, et al., Deep-learning-assisted diagnosis for knee magnetic resonance imaging: development and retrospective validation of MRNet, PLoS medicine 15 (11) (2018) e1002699.

[127] N. L. S. T. R. Team, Reduced lung-cancer mortality with low-dose computed tomographic screening, New England Journal of Medicine 365 (5) (2011) 395–409.

[128] N. J. Edwards, M. Oberti, R. R. Thangudu, S. Cai, P. B. McGarvey, S. Jacob, S. Madhavan, K. A. Ketchum, The CPTAC data portal: a resource for cancer proteomics research, Journal of proteome research 14 (6) (2015) 2707–2713.

[129] N. I. of Health, et al., National cancer institute. the cancer genome atlas program (2019).

[130] J. G. Elmore, G. M. Longton, P. A. Carney, B. M. Geller, T. Onega, A. N. Tosteson, H. D. Nelson, M. S. Pepe, K. H. Allison, S. J. Schnitt, et al., Diagnostic concordance among pathologists interpreting breast biopsy specimens, Jama 313 (11) (2015) 1122–1132.

[131] X. He, X. Yang, S. Zhang, J. Zhao, Y. Zhang, E. Xing, P. Xie, Sample-efficient deep learning for COVID-19 diagnosis based on CT scans, medrxiv (2020) 2020–04.

[132] K. Pogorelov, K. R. Randel, C. Griwodz, S. L. Eskeland, T. de Lange, D. Johansen, C. Spampinato, D.-T. Dang-Nguyen, M. Lux, P. T. Schmidt, et al., Kvasir: A multi-class image dataset for computer aided gastrointestinal disease detection, in: Proceedings of the 8th ACM on Multimedia Systems Conference, 2017, pp. 164–169.

[133] J. Yang, R. Shi, B. Ni, Medmnist classification decathlon: A lightweight automl benchmark for medical image analysis, in: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), IEEE, 2021, pp. 191–195.

[134] F. M. Bianchi, D. Grattarola, C. Alippi, Spectral clustering with graph neural networks for graph pooling, in: International Conference on Machine Learning, PMLR, 2020, pp. 874–883.

[135] Z. Zhong, L. Zheng, S. Li, Y. Yang, Generalizing a person retrieval model hetero-and homogeneously, in: Proceedings of the European conference on computer vision (ECCV), 2018, pp. 172–188.

[136] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, A. Torralba, Learning deep features for discriminative localization, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 2921–2929.

[137] S.-C. Huang, A. Pareek, M. Jensen, M. P. Lungren, S. Yeung, A. S. Chaudhari, Self-supervised learning for medical image classification: a systematic review and implementation guidelines, NPJ Digital Medicine 6 (1) (2023) 74.

[138] X. Chen, S. Xie, K. He, An empirical study of training self-supervised vision transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 9640–9649.

[139] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, G. E. Hinton, Big self-supervised models are strong semi-supervised learners, Advances in neural information processing systems 33 (2020) 22243–22255.

[140] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. Richemond, E. Buchatskaya, C. Doersch, B. Avila Pires, Z. Guo, M. Gheshlaghi Azar, et al., Bootstrap your own latent-a new approach to self-supervised learning, Advances in neural information processing systems 33 (2020) 21271–21284.

[141] V. Ljosa, K. L. Sokolnicki, A. E. Carpenter, Annotated high-throughput microscopy image sets for validation., Nature methods 9 (7) (2012) 637–637.

[142] Z. Peng, W. Huang, S. Gu, L. Xie, Y. Wang, J. Jiao, Q. Ye, Conformer: Local features coupling global representations for visual recognition, in: Proceedings of the IEEE/CVF international conference on computer vision, 2021, pp. 367–376.

[143] X. Chu, Z. Tian, Y. Wang, B. Zhang, H. Ren, X. Wei, H. Xia, C. Shen, Twins: Revisiting the design of spatial attention in vision transformers, Advances in Neural Information Processing Systems 34 (2021) 9355–9366.

[144] W. Yu, M. Luo, P. Zhou, C. Si, Y. Zhou, X. Wang, J. Feng, S. Yan, Metaformer is actually what you need for vision, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2022, pp. 10819–10829.

[145] H. Wu, B. Xiao, N. Codella, M. Liu, X. Dai, L. Yuan, L. Zhang, Cvt: Introducing convolutions to vision transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 22–31.

[146] X. Mao, G. Qi, Y. Chen, X. Li, R. Duan, S. Ye, Y. He, H. Xue, Towards robust vision transformer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 12042–12051.

[147] W. Xu, Y. Xu, T. Chang, Z. Tu, Co-scale conv-attentional image transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 9981–9990.

[148] B. Heo, S. Yun, D. Han, S. Chun, J. Choe, S. J. Oh, Rethinking spatial dimensions of vision transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 11936–11945.

[149] W. Wang, E. Xie, X. Li, D.-P. Fan, K. Song, D. Liang, T. Lu, P. Luo, L. Shao, Pyramid vision transformer: A versatile backbone for dense prediction without convolutions, in: Proceedings of the IEEE/CVF international conference on computer vision, 2021, pp. 568–578.

[150] W. H. Pinaya, P.-D. Tudosiu, J. Dafflon, P. F. Da Costa, V. Fernandez, P. Nachev, S. Ourselin, M. J. Cardoso, Brain imaging generation with latent diffusion models, in: MICCAI Workshop on Deep Generative Models, Springer, 2022, pp. 117–126.

[151] P. A. Moghadam, S. Van Dalen, K. C. Martin, J. Lennerz, S. Yip, H. Farahani, A. Bashashati, A morphology focused diffusion probabilistic model for synthesis of histopathology images, arXiv preprint arXiv:2209.13167 (2022).

[152] A. Kazerouni, E. K. Aghdam, M. Heidari, R. Azad, M. Fayyaz, I. Hacihaliloglu, D. Merhof, Diffusion models in medical imaging: A comprehensive survey, Medical Image Analysis (2023) 102846.

[153] H.-Y. Zhou, J. Guo, Y. Zhang, L. Yu, L. Wang, Y. Yu, nnformer: Interleaved transformer for volumetric segmentation, arXiv preprint arXiv:2109.03201 (2021).

[154] X. Huang, Z. Deng, D. Li, X. Yuan, Y. Fu, Missformer: An effective transformer for 2d medical image segmentation, IEEE Transactions on Medical Imaging (2022) 1–1 doi:10.1109/TMI.2022.3230943.

[155] S. Li, X. Sui, X. Luo, X. Xu, L. Yong, R. S. M. Goh, Medical image segmentation using squeeze-and-expansion transformers, in: The 30th International Joint Conference on Artificial Intelligence (IJCAI), 2021.

[156] W. Wang, C. Chen, M. Ding, H. Yu, S. Zha, J. Li, Transbts: Multimodal brain tumor segmentation using transformer, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 109–119.

[157] Y. Zhang, H. Liu, Q. Hu, Transfuse: Fusing transformers and cnns for medical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 14–24.

[158] J. M. J. Valanarasu, P. Oza, I. Hacihaliloglu, V. M. Patel, Medical transformer: Gated axial-attention for medical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 36–46.

[159] A. Hatamizadeh, Y. Tang, V. Nath, D. Yang, A. Myronenko, B. Landman, H. R. Roth, D. Xu, Unetr: Transformers for 3D medical image segmentation, in: Proceedings of the IEEE/CVF winter conference on applications of computer vision, 2022, pp. 574–584.

[160] A. Hatamizadeh, V. Nath, Y. Tang, D. Yang, H. R. Roth, D. Xu, Swin unetr: Swin transformers for semantic segmentation of brain tumors in MRI images, in: Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 7th International Workshop, BrainLes 2021, Held in Conjunction with MICCAI 2021, Virtual Event, September 27, 2021, Revised Selected Papers, Part I, Springer, 2022, pp. 272–284.

[161] Y. Tang, D. Yang, W. Li, H. R. Roth, B. Landman, D. Xu, V. Nath, A. Hatamizadeh, Self-supervised pre-training of swin transformers for 3D medical image analysis, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 20730–20740.

[162] Y. Xie, J. Zhang, C. Shen, Y. Xia, CoTr: Efficiently bridging CNN and transformer for 3D medical image segmentation, arXiv preprint arXiv:2103.03024 (2021).

[163] D. Yang, A. Myronenko, X. Wang, Z. Xu, H. R. Roth, D. Xu, T-AutoML: Automated machine learning for lesion segmentation using transformers in 3D medical imaging, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 3962–3974.

[164] X. Luo, M. Hu, T. Song, G. Wang, S. Zhang, Semi-supervised medical image segmentation via cross teaching between CNN and transformer, in: International Conference on Medical Imaging with Deep Learning, PMLR, 2022, pp. 820–833.

[165] L. Zhou, H. Liu, J. Bae, J. He, D. Samaras, P. Prasanna, Self pre-training with masked autoencoders for medical image analysis, arXiv preprint arXiv:2203.05573 (2022).

[166] R. Azad, E. K. Aghdam, A. Rauland, Y. Jia, A. H. Avval, A. Bozorgpour, S. Karimijafarbigloo, J. P. Cohen, E. Adeli, D. Merhof, Medical image segmentation review: The success of U-Net, arXiv preprint arXiv:2211.14830 (2022).

[167] M. Asadi-Aghbolaghi, R. Azad, M. Fathy, S. Escalera, Multi-level context gating of embedded collective knowledge for medical image segmentation, arXiv preprint arXiv:2003.05056 (2020).

[168] S. Karimijafarbigloo, R. Azad, A. Kazerouni, S. Ebadollahi, D. Merhof, MMCFormer: Missing modality compensation transformer for brain tumor segmentation, in: Medical Imaging with Deep Learning, 2023.

[169] E. K. Aghdam, R. Azad, M. Zarvani, D. Merhof, Attention swin u-net: Cross-contextual attention mechanism for skin lesion segmentation, arXiv preprint arXiv:2210.16898 (2022).

[170] B. Landman, Z. Xu, J. Igelsias, M. Styner, T. Langerak, A. Klein, Miccai multi-atlas labeling beyond the cranial vault–workshop and challenge, in: Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge, Vol. 5, 2015, p. 12.

[171] M. Nolden, S. Zelzer, A. Seitel, D. Wald, M. Müller, A. M. Franz, D. Maleike, M. Fangerau, M. Baumhauer, L. Maier-Hein, et al., The medical imaging interaction toolkit: challenges and advances, International journal of computer assisted radiology and surgery 8 (4) (2013) 607–620.

[172] R. Azad, M. T. Al-Antary, M. Heidari, D. Merhof, Transnorm: Transformer provides a strong spatial normalization mechanism for a deep segmentation model, IEEE Access 10 (2022) 108205–108215.

[173] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, B. Glocker, D. Rueckert, Attention U-Net: Learning where to look for the pancreas, in: Medical Imaging with Deep Learning, 2018.
URL https://openreview.net/forum?id=Skft7cijM

[174] F. Isensee, P. F. Jäger, P. M. Full, P. Vollmuth, K. H. Maier-Hein, nnU-net for brain tumor segmentation, in: International MICCAI Brainlesion Workshop, Springer, 2020, pp. 118–132.

[175] F. Liu, X. Ren, Z. Zhang, X. Sun, Y. Zou, Rethinking skip connection with layer normalization in transformers and resnets, arXiv preprint arXiv:2105.07205 (2021).

[176] L.-C. Chen, G. Papandreou, F. Schroff, H. Adam, Rethinking atrous convolution for semantic image segmentation, arXiv preprint arXiv:1706.05587 (2017).

[177] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz, et al., Attention u-net: Learning where to look for the pancreas, arXiv preprint arXiv:1804.03999 (2018).

[178] S. Zheng, J. Lu, H. Zhao, X. Zhu, Z. Luo, Y. Wang, Y. Fu, J. Feng, T. Xiang, P. H. 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, 2021, pp. 6881–6890.

[179] S. Woo, J. Park, J.-Y. Lee, I. S. Kweon, Cbam: Convolutional block attention module, in: Proceedings of the European conference on computer vision (ECCV), 2018, pp. 3–19.

[180] J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, D. Rueckert, Attention gated networks: Learning to leverage salient regions in medical images, Medical image analysis 53 (2019) 197–207.

[181] H. Wang, Y. Zhu, B. Green, H. Adam, A. Yuille, L.-C. Chen, Axial-deeplab: Stand-alone axial-attention for panoptic segmentation, in: European Conference on Computer Vision, Springer, 2020, pp. 108–126.

[182] F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, K. H. Maier-Hein, nnu-net: a self-configuring method for deep learning-based biomedical image segmentation, Nature Methods 18 (2) (2021) 203–+.
URL <GotoISI>://WOS:000599000100001

[183] Y. Tang, R. Gao, H. H. Lee, S. Han, Y. Chen, D. Gao, V. Nath, C. Bermudez, M. R. Savona, R. G. Abramson, et al., High-resolution 3d abdominal segmentation with random patch network fusion, Medical Image Analysis 69 (2021) 101894.

[184] Y. Zhou, Z. Li, S. Bai, C. Wang, X. Chen, M. Han, E. Fishman, A. L. Yuille, Prior-aware neural network for partially-supervised multi-organ segmentation, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 10672–10681.

[185] A. Myronenko, 3D MRI brain tumor segmentation using autoencoder regularization, in: International MICCAI Brainlesion Workshop, Springer, 2018, pp. 311–320.

[186] S. Liu, L. Qi, H. Qin, J. Shi, J. Jia, Path aggregation network for instance segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 8759–8768.

[187] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, S. Zagoruyko, End-to-end object detection with transformers, in: European conference on computer vision, Springer, 2020, pp. 213–229.

[188] J. Lee, Y. Lee, J. Kim, A. Kosiorek, S. Choi, Y. W. Teh, Set transformer: A framework for attention-based permutation-invariant neural networks, in: International Conference on Machine Learning, PMLR, 2019, pp. 3744–3753.

[189] A. Gupta, S. Gehlot, S. Goswami, S. Motwani, R. Gupta, Á. G. Faura, D. Štepec, T. Martinčič, R. Azad, D. Merhof, et al., SegPC-2021: A challenge & dataset on segmentation of multiple myeloma plasma cells from microscopic images, Medical Image Analysis 83 (2023) 102677.

[190] X. Du, T.-Y. Lin, P. Jin, G. Ghiasi, M. Tan, Y. Cui, Q. V. Le, X. Song, Spinenet: Learning scale-permuted backbone for recognition and localization, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 11592–11601.

[191] C. Liu, L.-C. Chen, F. Schroff, H. Adam, W. Hua, A. L. Yuille, L. Fei-Fei, Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 82–92.

[192] W. Bae, S. Lee, Y. Lee, B. Park, M. Chung, K.-H. Jung, Resource optimized neural architecture search for 3D medical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2019, pp. 228–236.

[193] S. Kim, I. Kim, S. Lim, W. Baek, C. Kim, H. Cho, B. Yoon, T. Kim, Scalable neural architecture search for 3D medical image segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2019, pp. 220–228.

[194] S. Qiao, W. Shen, Z. Zhang, B. Wang, A. Yuille, Deep co-training for semi-supervised image recognition, in: Proceedings of the european conference on computer vision (eccv), 2018, pp. 135–152.

[195] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, M. Sugiyama, Co-teaching: Robust training of deep neural networks with extremely noisy labels, Advances in neural information processing systems 31 (2018).

[196] X. Chen, Y. Yuan, G. Zeng, J. Wang, Semi-supervised semantic segmentation with cross pseudo supervision, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 2613–2622.

[197] L. Yu, S. Wang, X. Li, C.-W. Fu, P.-A. Heng, Uncertainty-aware self-ensembling model for semi-supervised 3D left atrium segmentation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2019, pp. 605–613.

[198] O. Bernard, A. Lalande, C. Zotti, F. Cervenansky, X. Yang, P.-A. Heng, I. Cetin, K. Lekadir, O. Camara, M. A. G. 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) (2018) 2514–2525.

[199] T. Mendonça, P. M. Ferreira, J. S. Marques, A. R. S. Marcal, J. Rozeira, Ph2 - a dermoscopic image database for research and benchmarking, in: 2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), 2013, pp. 5437–5440. doi:10.1109/EMBC.2013.6610779.

[200] B. H. Menze, A. Jakab, S. Bauer, J. Kalpathy-Cramer, K. Farahani, J. Kirby, Y. Burren, N. Porz, J. Slotboom, R. Wiest, et al., The multimodal brain tumor image segmentation benchmark (BRATS), IEEE transactions on medical imaging 34 (10) (2014) 1993–2024.

[201] S. Bakas, H. Akbari, A. Sotiras, M. Bilello, M. Rozycki, J. S. Kirby, J. B. Freymann, K. Farahani, C. Davatzikos, Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features, Scientific data 4 (1) (2017) 1–13.

[202] S. Bakas, M. Reyes, A. Jakab, S. Bauer, M. Rempfler, A. Crimi, R. T. Shinohara, C. Berger, S. M. Ha, M. 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).

[203] K. Sirinukunwattana, J. P. Pluim, H. Chen, X. Qi, P.-A. Heng, Y. B. Guo, L. Y. Wang, B. J. Matuszewski, E. Bruni, U. Sanchez, et al., Gland segmentation in colon histology images: The glas challenge contest, Medical image analysis 35 (2017) 489–502.

[204] N. Kumar, R. Verma, S. Sharma, S. Bhargava, A. Vahadane, A. Sethi, A dataset and a technique for generalized nuclear segmentation for computational pathology, IEEE transactions on medical imaging 36 (7) (2017) 1550–1560.

[205] A. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B. Landman, G. Litjens, B. Menze, et al., A large annotated medical image dataset for the development and evaluation of segmentation algorithms. arxiv 2019, arXiv preprint arXiv:1902.09063.

[206] U. Baid, S. Ghodasara, S. Mohan, M. Bilello, E. Calabrese, E. Colak, K. Farahani, J. Kalpathy-Cramer, F. C. Kitamura, S. Pati, et al., The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification, arXiv preprint arXiv:2107.02314 (2021).

[207] A. Gupta, P. Mallick, O. Sharma, R. Gupta, R. Duggal, PCSeg: Color model driven probabilistic multiphase level set based tool for plasma cell segmentation in multiple myeloma, PloS one 13 (12) (2018) e0207908.

[208] A. Gupta, R. Duggal, S. Gehlot, R. Gupta, A. Mangal, L. Kumar, N. Thakkar, D. Satpathy, GCTI-SN: Geometry-inspired chemical and tissue invariant stain normalization of microscopic medical images, Medical Image Analysis 65 (2020) 101788.

[209] S. Gehlot, A. Gupta, R. Gupta, Ednfc-net: Convolutional neural network with nested feature concatenation for nuclei-instance segmentation, in: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2020, pp. 1389–1393.

[210] J. I. Orlando, H. Fu, J. B. Breda, K. van Keer, D. R. Bathula, A. Diaz-Pinto, R. Fang, P.-A. Heng, J. Kim, J. Lee, et al., Refuge challenge: A unified framework

for evaluating automated methods for glaucoma assessment from fundus photographs, Medical image analysis 59 (2020) 101570.

[211] D.-P. Fan, G.-P. Ji, T. Zhou, G. Chen, H. Fu, J. Shen, L. Shao, Pranet: Parallel reverse attention network for polyp segmentation, in: International conference on medical image computing and computer-assisted intervention, Springer, 2020, pp. 263–273.

[212] D. Jha, P. H. Smedsrud, M. A. Riegler, P. Halvorsen, T. d. Lange, D. Johansen, H. D. Johansen, Kvasir-seg: A segmented polyp dataset, in: International Conference on Multimedia Modeling, Springer, 2020, pp. 451–462.

[213] A. L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B. A. Landman, G. Litjens, B. Menze, et al., A large annotated medical image dataset for the development and evaluation of segmentation algorithms, arXiv preprint arXiv:1902.09063 (2019).

[214] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, A. 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) (2017) 834–848.

[215] H. Bao, L. Dong, S. Piao, F. Wei, BEit: BERT pre-training of image transformers, in: International Conference on Learning Representations, 2022.
URL https://openreview.net/forum?id=p-BhZSz59o4

[216] Z. Xie, Z. Zhang, Y. Cao, Y. Lin, J. Bao, Z. Yao, Q. Dai, H. Hu, Simmim: A simple framework for masked image modeling, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 9653–9663.

[217] A. El-Nouby, G. Izacard, H. Touvron, I. Laptev, H. Jegou, E. Grave, Are large-scale datasets necessary for self-supervised pre-training?, arXiv preprint arXiv:2112.10740 (2021).

[218] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, R. Girshick, Masked autoencoders are scalable vision learners, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 16000–16009.

[219] M. M. Rahman, R. Marculescu, Medical image segmentation via cascaded attention decoding, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2023, pp. 6222–6231.

[220] R. Azad, R. Arimond, E. K. Aghdam, A. Kazerouni, D. Merhof, DAE-Former: Dual attention-guided efficient transformer for medical image segmentation, arXiv preprint arXiv:2212.13504 (2022).

[221] H. Huang, S. Xie, L. Lin, Y. Iwamoto, X. Han, Y.-W. Chen, R. Tong, Scaleformer: Revisiting the transformer-based backbones from a scale-wise perspective for medical image segmentation, arXiv preprint arXiv:2207.14552 (2022).

[222] R. Azad, L. Niggemeier, M. Huttemann, A. Kazerouni, E. K. Aghdam, Y. Velichko, U. Bagci, D. Merhof, Beyond self-attention: Deformable large kernel attention for medical image segmentation (2023). arXiv:2309.00121.

[223] M. M. Rahman, R. Marculescu, Multi-scale hierarchical vision transformer with cascaded attention decoding for medical image segmentation, in: Medical Imaging with Deep Learning, 2023.
URL https://openreview.net/forum?id=u0MHV19E2n

[224] T. Mendonça, P. M. Ferreira, J. S. Marques, A. R. Marcal, J. 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), IEEE, 2013, pp. 5437–5440.

[225] Z. Zhang, L. Yu, X. Liang, W. Zhao, L. Xing, TransCT: dual-path transformer for low dose computed tomography, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24, Springer, 2021, pp. 55–64.

[226] D. Wang, Z. Wu, H. Yu, TED-net: Convolution-free T2T vision transformer-based encoder-decoder dilation network for low-dose CT denoising, in: Machine Learning in Medical Imaging: 12th International Workshop, MLMI 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 12, Springer, 2021, pp. 416–425.

[227] A. Luthra, H. Sulakhe, T. Mittal, A. Iyer, S. Yadav, Eformer: Edge enhancement based transformer for medical image denoising, arXiv preprint arXiv:2109.08044 (2021).

[228] Y. Luo, Y. Wang, C. Zu, B. Zhan, X. Wu, J. Zhou, D. Shen, L. Zhou, 3D transformer-GAN for high-quality PET reconstruction, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 276–285.

[229] L. Zhang, Z. Xiao, C. Zhou, J. Yuan, Q. He, Y. Yang, X. Liu, D. Liang, H. Zheng, W. Fan, et al., Spatial adaptive and transformer fusion network (STFNet) for low-count PET blind denoising with MRI, Medical Physics 49 (1) (2022) 343–356.

[230] D. Wang, F. Fan, Z. Wu, R. Liu, F. Wang, H. Yu, CTformer: convolution-free token2token dilated vision transformer for low-dose CT denoising, Physics in Medicine & Biology 68 (6) (2023) 065012.

[231] L. Yang, D. Zhang, et al., Low-dose CT denoising via sinogram inner-structure transformer, arXiv preprint arXiv:2204.03163 (2022).

[232] C. Wang, K. Shang, H. Zhang, Q. Li, Y. Hui, S. K. Zhou, Dudotrans: Dual-domain transformer provides more attention for sinogram restoration in sparse-view CT reconstruction, arXiv preprint arXiv:2111.10790 (2021).

[233] T.-O. Buchholz, F. Jug, Fourier image transformer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 1846–1854.

[234] C. Shi, Y. Xiao, Z. Chen, Dual-domain sparse-view CT reconstruction with transformers, Physica Medica 101 (2022) 1–7.

[235] M. Wu, Y. Xu, Y. Xu, G. Wu, Q. Chen, H. Lin, Adaptively re-weighting multi-loss untrained transformer for sparse-view cone-beam CT reconstruction, arXiv preprint arXiv:2203.12476 (2022).

[236] K. Lin, R. Heckel, Vision transformers enable fast and robust accelerated MRI, in: Medical Imaging with Deep Learning, 2021.

[237] C.-M. Feng, Y. Yan, H. Fu, L. Chen, Y. Xu, Task transformer network for joint MRI reconstruction and super-resolution, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 307–317.

[238] D. Mahapatra, Z. Ge, MR image super resolution by combining feature disentanglement CNNs and vision transformers, in: Medical Imaging with Deep Learning, 2021.

[239] C. Fang, D. Zhang, L. Wang, Y. Zhang, L. Cheng, J. Han, Cross-modality high-frequency transformer for MR image super-resolution, in: Proceedings of the 30th ACM International Conference on Multimedia, 2022, pp. 1584–1592.

[240] Y. Korkmaz, S. U. Dar, M. Yurt, M. Özbey, T. Cukur, Unsupervised MRI reconstruction via zero-shot learned adversarial transformers, IEEE Transactions on Medical Imaging 41 (7) (2022) 1747–1763.

[241] B. Zhou, N. Dey, J. Schlemper, S. S. M. Salehi, C. Liu, J. S. Duncan, M. Sofka, DSFormer: a dual-domain self-supervised transformer for accelerated multi-contrast MRI reconstruction, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2023, pp. 4966–4975.

[242] E. Seeram, Computed Tomography-E-Book: Physical Principles, Clinical Applications, and Quality Control, Elsevier Health Sciences, 2015.

[243] J. P. Mathews, Q. P. Campbell, H. Xu, P. Halleck, A review of the application of X-ray computed tomography to the study of coal, Fuel 209 (2017) 10–24.

[244] D. J. Brenner, E. J. Hall, Computed tomography—an increasing source of radiation exposure, New England journal of medicine 357 (22) (2007) 2277–2284.

[245] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, J. K. Seo, Deep learning for undersampled MRI reconstruction, Physics in Medicine & Biology 63 (13) (2018) 135007.

[246] H.-M. Zhang, B. Dong, A review on deep learning in medical image reconstruction, Journal of the Operations Research Society of China 8 (2) (2020) 311–340.

[247] L. Yuan, Y. Chen, T. Wang, W. Yu, Y. Shi, Z.-H. Jiang, F. E. Tay, J. Feng, S. Yan, Tokens-to-token vit: Training vision transformers from scratch on imagenet, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 558–567.

[248] Z. Wang, X. Cun, J. Bao, W. Zhou, J. Liu, H. Li, Uformer: A general u-shaped transformer for image restoration, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 17683–17693.

[249] A. Buades, B. Coll, J.-M. Morel, A non-local algorithm for image denoising, in: 2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05), Vol. 2, Ieee, 2005, pp. 60–65.

[250] K. Dabov, A. Foi, V. Katkovnik, K. Egiazarian, Image denoising with block-matching and 3D filtering, in: Image processing: algorithms and systems, neural networks, and machine learning, Vol. 6064, SPIE, 2006, pp. 354–365.

[251] C. Wang, Z. Hu, P. Shi, H. Liu, Low dose PET reconstruction with total variation regularization, in: 2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, IEEE, 2014, pp. 1917–1920.

[252] C. H. McCollough, A. C. Bartley, R. E. Carter, B. Chen, T. A. Drees, P. Edwards, D. R. Holmes III, A. E. Huang, F. Khan, S. 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) (2017) e339–e352.

[253] J. Bian, J. H. Siewerdsen, X. Han, E. Y. Sidky, J. L. Prince, C. A. Pelizzari, X. Pan, Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT, Physics in Medicine & Biology 55 (22) (2010) 6575.

[254] Y. Han, J. C. Ye, Framing U-Net via deep convolutional framelets: Application to sparse-view ct, IEEE transactions on medical imaging 37 (6) (2018) 1418–1429.

[255] A. C. Kak, M. Slaney, Principles of computerized tomographic imaging, SIAM, 2001.

[256] X. Pan, E. Y. Sidky, M. Vannier, Why do commercial CT scanners still employ traditional, filtered backprojection for image reconstruction?, Inverse problems 25 (12) (2009) 123009.

[257] L. A. Feldkamp, L. C. Davis, J. W. Kress, Practical cone-beam algorithm, Josa a 1 (6) (1984) 612–619.

[258] S. Patel, J. Brown, T. Pimentel, R. Kelly, F. Abella, C. Durack, Cone beam computed tomography in endodontics–a review of the literature, International endodontic journal 52 (8) (2019) 1138–1152.

[259] D. Ulyanov, A. Vedaldi, V. Lempitsky, Deep image prior, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 9446–9454.

[260] J. Johnson, A. Alahi, L. Fei-Fei, Perceptual losses for real-time style transfer and super-resolution, in: European conference on computer vision, Springer, 2016, pp. 694–711.

[261] J. Leuschner, M. Schmidt, D. O. Baguer, P. Maaß, The lodopab-ct dataset: A benchmark dataset for low-dose CT reconstruction methods, arXiv preprint arXiv:1910.01113 (2019).

[262] J. Zbontar, F. Knoll, A. Sriram, T. Murrell, Z. Huang, M. J. Muckley, A. Defazio, R. Stern, P. Johnson, M. Bruno, M. Parente, K. J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdzal, A. Romero, M. Rabbat, P. Vincent, N. Yakubova, J. Pinkerton, D. Wang, E. Owens, C. L. Zitnick, M. P. Recht, D. K. Sodickson, Y. W. Lui, fastMRI: An open dataset and benchmarks for accelerated MRI, 2018. arXiv:1811.08839.

[263] T. R. Moen, B. Chen, D. R. Holmes III, X. Duan, Z. Yu, L. Yu, S. Leng, J. G. Fletcher, C. H. McCollough, Low-dose CT image and projection dataset, Medical physics 48 (2) (2021) 902–911.

[264] S. G. Armato III, G. McLennan, L. Bidaut, M. F. McNitt-Gray, C. R. Meyer, A. P. Reeves, B. Zhao, D. R. Aberle, C. I. Henschke, E. A. Hoffman, et al., The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans, Medical physics 38 (2) (2011) 915–931.

[265] B. I. A. Group, Ixi dataset, http://brain-development.org/ixi-dataset/.

[266] C.-C. Shieh, Y. Gonzalez, B. Li, X. Jia, S. Rit, C. Mory, M. Riblett, G. Hugo, Y. Zhang, Z. Jiang, et al., SPARE: Sparse-view reconstruction challenge for 4D cone-beam CT from a 1-min scan, Medical physics 46 (9) (2019) 3799–3811.

[267] H. Der Sarkissian, F. Lucka, M. van Eijnatten, G. Colacicco, S. B. Coban, K. J. Batenburg, A cone-beam X-ray computed tomography data collection designed for machine learning, Scientific data 6 (1) (2019) 1–8.

[268] W.-A. Lin, H. Liao, C. Peng, X. Sun, J. Zhang, J. Luo, R. Chellappa, S. K. Zhou, Dudonet: Dual domain network for ct metal artifact reduction, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 10512–10521.

[269] E. Plenge, D. H. Poot, M. Bernsen, G. Kotek, G. Houston, P. Wielopolski, L. van der Weerd, W. J. Niessen, E. Meijering, Super-resolution methods in MRI: can they improve the trade-off between resolution, signal-to-noise ratio, and acquisition time?, Magnetic resonance in medicine 68 (6) (2012) 1983–1993.

[270] S. d'Ascoli, H. Touvron, M. L. Leavitt, A. S. Morcos, G. Biroli, L. Sagun, Convit: Improving vision transformers with soft convolutional inductive biases, in: International Conference on Machine Learning, PMLR, 2021, pp. 2286–2296.

[271] B. Lim, S. Son, H. Kim, S. Nah, K. Mu Lee, Enhanced deep residual networks for single image super-resolution, in: Proceedings of the IEEE conference on computer vision and pattern recognition workshops, 2017, pp. 136–144.

[272] C.-M. Feng, Y. Yan, H. Fu, L. Chen, Y. Xu, Task transformer network for joint MRI reconstruction and super-resolution, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24, Springer, 2021, pp. 307–317.

[273] C.-M. Feng, Y. Yan, G. Chen, Y. Xu, Y. Hu, L. Shao, H. Fu, Multi-modal transformer for accelerated MR imaging, IEEE Transactions on Medical Imaging (2022).

[274] X. Zhang, X. He, J. Guo, N. Ettehadi, N. Aw, D. Semanek, J. Posner, A. Laine, Y. Wang, PTNet: a high-resolution infant MRI synthesizer based on transformer, arXiv preprint arXiv:2105.13993 (2021).

[275] O. Dalmaz, M. Yurt, T. Çukur, ResViT: Residual vision transformers for multimodal medical image synthesis, IEEE Transactions on Medical Imaging 41 (10) (2022) 2598–2614. doi:10.1109/TMI.2022.3167808.

[276] J. Liu, S. Pasumarthi, B. Duffy, E. Gong, G. Zaharchuk, K. Datta, One model to synthesize them all: Multi-contrast multi-scale transformer for missing data imputation, arXiv preprint arXiv:2204.13738 (2022).

[277] N.-C. Ristea, A.-I. Miron, O. Savencu, M.-I. Georgescu, N. Verga, F. S. Khan, R. T. Ionescu, CyTran: Cycle-consistent transformers for non-contrast to contrast CT translation, arXiv preprint arXiv:2110.06400 (2021).

[278] S. A. Kamran, K. F. Hossain, A. Tavakkoli, S. L. Zucker-brod, S. A. Baker, Vtgan: Semi-supervised retinal image synthesis and disease prediction using vision transformers, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 3235–3245.

[279] Y. Li, T. Zhou, K. He, Y. Zhou, D. Shen, SLMT-Net: A self-supervised learning based multi-scale transformer network for cross-modality mr image synthesis, arXiv preprint arXiv:2212.01108 (2022).

[280] K. M. Choromanski, V. Likhosherstov, D. Dohan, X. Song, A. Gane, T. Sarlos, P. Hawkins, J. Q. Davis, A. Mohiuddin, L. Kaiser, D. B. Belanger, L. J. Colwell, A. Weller, Rethinking attention with performers, in: International Conference on Learning Representations, 2021.
URL https://openreview.net/forum?id=Ua6zuk0WRH

[281] A. Makropoulos, E. C. Robinson, A. Schuh, R. Wright, S. Fitzgibbon, J. Bozek, S. J. Counsell, J. Steinweg, K. Vecchiato, J. Passerat-Palmbach, et al., The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction, Neuroimage 173 (2018) 88–112.

[282] P. Isola, J.-Y. Zhu, T. Zhou, A. A. Efros, Image-to-image translation with conditional adversarial networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 1125–1134.

[283] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, B. Catanzaro, High-resolution image synthesis and semantic manipulation with conditional gans, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 8798–8807.

[284] S. U. Dar, M. Yurt, L. Karacan, A. Erdem, E. Erdem, T. Cukur, Image synthesis in multi-contrast MRI with conditional generative adversarial networks, IEEE transactions on medical imaging 38 (10) (2019) 2375–2388.

[285] H. Zhang, I. Goodfellow, D. Metaxas, A. Odena, Self-attention generative adversarial networks, in: International conference on machine learning, PMLR, 2019, pp. 7354–7363.

[286] T. Nyholm, S. Svensson, S. Andersson, J. Jonsson, M. Sohlin, C. Gustafsson, E. Kjellén, K. Söderström, P. Albertsson, L. 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) (2018) 1295–1300.

[287] O. Maier, B. H. Menze, J. von der Gablentz, L. Häni, M. P. Heinrich, M. Liebrand, S. Winzeck, A. Basit, P. Bentley, L. Chen, et al., ISLES 2015-A public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI, Medical Image Analysis 35 (2017) 250–269.

[288] S. Hajeb Mohammad Alipour, H. Rabbani, M. R. Akhlaghi, Diabetic retinopathy grading by digital curvelet transform, Computational and mathematical methods in medicine 2012 (2012).

[289] X. Ma, G. Luo, W. Wang, K. Wang, Transformer network for significant stenosis detection in ccta of coronary arteries, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24, Springer, 2021, pp. 516–525.

[290] H. Jiang, P. Zhang, C. Che, B. Jin, RDFNet: A fast caries detection method incorporating transformer mechanism, Computational and Mathematical Methods in Medicine 2021 (2021).

[291] R. Wagner, K. Rohr, Cellcentroidformer: Combining self-attention and convolution for cell detection, in: Annual Conference on Medical Image Understanding and Analysis, Springer, 2022, pp. 212–222.

[292] D. M. Nguyen, H. Nguyen, M. T. Truong, T. Cao, B. T. Nguyen, N. Ho, P. Swoboda, S. Albarqouni, P. Xie, D. Sonntag, Joint self-supervised image-volume representation learning with intra-inter contrastive clustering, in: 37th AAAI Conference on Artificial Intelligence, AAAI, 2022.

[293] Z. Shen, R. Fu, C. Lin, S. Zheng, COTR: Convolution in transformer network for end to end polyp detection, in: 2021 7th International Conference on Computer and Communications (ICCC), IEEE, 2021, pp. 1757–1761.

[294] Q. Kong, Y. Wu, C. Yuan, Y. Wang, CT-CAD: Context-aware transformers for end-to-end chest abnormality detection on X-Rays, in: 2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), IEEE, 2021, pp. 1385–1388.

[295] R. Tao, G. Zheng, Spine-transformers: Vertebra detection and localization in arbitrary field-of-view spine ct with transformers, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 93–103.

[296] B. Wittmann, F. Navarro, S. Shit, B. Menze, et al., Focused decoding enables 3D anatomical detection by transformers, Machine Learning for Biomedical Imaging 2 (February 2023 issue) (2023) 72–95.

[297] Z. Yao, J. Ai, B. Li, C. Zhang, Efficient detr: improving end-to-end object detector with dense prior, arXiv preprint arXiv:2104.01318 (2021).

[298] A. Criminisi, J. Shotton, S. Bucciarelli, Decision forests with long-range spatial context for organ localization in CT volumes, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), Citeseer, 2009, pp. 69–80.

[299] F. Li, H. Zhang, S. Liu, J. Guo, L. M. Ni, L. Zhang, Dn-detr: Accelerate detr training by introducing query denoising, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 13619–13627.

[300] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. Ni, H.-Y. Shum, DINO: DETR with improved denoising anchor boxes for end-to-end object detection, in: The Eleventh International Conference on Learning Representations, 2023.
URL https://openreview.net/forum?id=3mRwyG5one

[301] T. Prangemeier, C. Reich, H. Koeppl, Attention-based transformers for instance segmentation of cells in microstructures, in: 2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), IEEE, 2020, pp. 700–707.

[302] K. He, X. Zhang, S. Ren, J. Sun, Spatial pyramid pooling in deep convolutional networks for visual recognition, IEEE transactions on pattern analysis and machine intelligence 37 (9) (2015) 1904–1916.

[303] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, S. Belongie, Feature pyramid networks for object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 2117–2125.

[304] S. Liu, L. Qi, H. Qin, J. Shi, J. Jia, Path aggregation network for instance segmentation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 8759–8768.

[305] M. Tan, Q. Le, Efficientnetv2: Smaller models and faster training, in: International Conference on Machine Learning, PMLR, 2021, pp. 10096–10106.

[306] M. Raghu, T. Unterthiner, S. Kornblith, C. Zhang, A. Dosovitskiy, Do vision transformers see like convolutional neural networks?, Advances in Neural Information Processing Systems 34 (2021) 12116–12128.

[307] O. Schoppe, C. Pan, J. Coronel, H. Mai, Z. Rong, M. I. Todorov, A. Müskes, F. Navarro, H. Li, A. Ertürk, et al., Deep learning-enabled multi-organ segmentation in whole-body mouse scans, Nature communications 11 (1) (2020) 1–14.

[308] K. H. Hohne, M. Bomans, M. Riemer, R. Schubert, U. Tiede, W. Lierse, A volume-based anatomical atlas, IEEE Computer Graphics and Applications 12 (04) (1992) 73–77.

[309] V. Ulman, M. Maška, K. E. Magnusson, O. Ronneberger, C. Haubold, N. Harder, P. Matula, P. Matula, D. Svoboda, M. Radojevic, et al., An objective comparison of cell-tracking algorithms, Nature methods 14 (12) (2017) 1141–1152.

[310] O. Jimenez-del Toro, H. Müller, M. Krenn, K. Gruenberg, A. A. Taha, M. Winterstein, I. Eggel, A. Foncubierta-Rodríguez, O. Goksel, A. Jakab, et al., Cloud-based evaluation of anatomical structure segmentation and landmark detection algorithms: VISCERAL anatomy benchmarks, IEEE transactions on medical imaging 35 (11) (2016) 2459–2475.

[311] Y. Ji, H. Bai, C. Ge, J. Yang, Y. Zhu, R. Zhang, Z. Li, L. Zhanng, W. Ma, X. Wan, et al., Amos: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation, Advances in Neural Information Processing Systems 35 (2022) 36722–36732.

[312] J. Bernal, F. J. Sánchez, G. Fernández-Esparrach, D. Gil, C. Rodríguez, F. Vilariño, WM-DOVA maps for accurate polyp highlighting in colonoscopy: Validation vs. saliency maps from physicians, Computerized medical imaging and graphics 43 (2015) 99–111.

[313] J. Silva, A. Histace, O. Romain, X. Dray, B. 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) (2014) 283–293.

[314] J. Bernal, J. Sánchez, F. Vilarino, Towards automatic polyp detection with a polyp appearance model, Pattern Recognition 45 (9) (2012) 3166–3182.

[315] N. T. Nguyen, P. T. Truong, V. T. Ho, T. V. Nguyen, H. T. Pham, M. T. Nguyen, L. T. Dam, H. Q. Nguyen, VinDr Lab: A data platform for medical AI, URL: https://github. com/vinbigdata-medical/vindr-lab (2021).

[316] J. Liu, J. Lian, Y. Yu, ChestX-det10: chest X-ray dataset on detection of thoracic abnormalities, arXiv preprint arXiv:2006.10550 (2020).

[317] A. Sekuboyina, A. Bayat, M. E. Husseini, M. Löffler, M. Rempfler, J. Kukačka, G. Tetteh, A. Valentinitsch, C. Payer, M. Urschler, et al., Verse: a vertebrae labelling and segmentation benchmark, arXiv. org e-Print archive (2020).

[318] B. Glocker, D. Zikic, E. Konukoglu, D. R. Haynor, A. Criminisi, Vertebrae localization in pathological spine CT via dense classification from sparse annotations, in: International conference on medical image computing and computer-assisted intervention, Springer, 2013, pp. 262–270.

[319] H. A. Qadir, I. Balasingham, J. Solhusvik, J. Bergsland, L. Aabakken, Y. Shin, Improving automatic polyp detection using CNN by exploiting temporal dependency in colonoscopy video, IEEE journal of biomedical and health informatics 24 (1) (2019) 180–193.

[320] H. A. Qadir, Y. Shin, J. Solhusvik, J. Bergsland, L. Aabakken, I. Balasingham, Toward real-time polyp detection using fully CNNs for 2D gaussian shapes prediction, Medical Image Analysis 68 (2021) 101897.

[321] Z. Cai, N. Vasconcelos, Cascade r-cnn: Delving into high quality object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 6154–6162.

[322] J. Redmon, S. Divvala, R. Girshick, A. Farhadi, You only look once: Unified, real-time object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 779–788.

[323] H. Rezatofighi, N. Tsoi, J. Gwak, A. Sadeghian, I. Reid, S. Savarese, Generalized intersection over union: A metric and a loss for bounding box regression, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 658–666.

[324] S. Qiao, L.-C. Chen, A. Yuille, Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2021, pp. 10213–10224.

[325] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, S. Belongie, Feature pyramid networks for object detection, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 2117–2125.

[326] Q. Chen, Y. Wang, T. Yang, X. Zhang, J. Cheng, J. Sun, You only look one-level feature, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2021, pp. 13039–13048.

[327] Z. Zheng, P. Wang, W. Liu, J. Li, R. Ye, D. Ren, Distance-iou loss: Faster and better learning for bounding box regression, in: Proceedings of the AAAI conference on artificial intelligence, Vol. 34, 2020, pp. 12993–13000.

[328] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, O. Ronneberger, 3D U-Net: learning dense volumetric segmentation from sparse annotation, in: International conference on medical image computing and computer-assisted intervention, Springer, 2016, pp. 424–432.

[329] G. Haskins, U. Kruger, P. Yan, Deep learning in medical image registration: a survey, Machine Vision and Applications 31 (1) (2020) 1–18.

[330] F. Alam, S. U. Rahman, S. Ullah, K. Gulati, Medical image registration in image guided surgery: Issues, challenges and research opportunities, Biocybernetics and Biomedical Engineering 38 (1) (2018) 71–89.

[331] F. Alam, S. U. Rahman, Challenges and solutions in multimodal medical image subregion detection and registration, Journal of medical imaging and radiation sciences 50 (1) (2019) 24–30.

[332] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, A. V. Dalca, An unsupervised learning model for deformable medical image registration, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2018, pp. 9252–9260.

[333] J. Chen, Y. Li, Y. Du, E. C. Frey, Generating anthropomorphic phantoms using fully unsupervised deformable image registration with convolutional neural networks, Medical physics 47 (12) (2020) 6366–6380.

[334] L. Chen, Q. T. Yu, Transformers make strong encoders for medical image segmentation. arxiv 2021, arXiv preprint arXiv:2102.04306.

[335] T. Lin, Y. Wang, X. Liu, X. Qiu, A survey of transformers, arXiv preprint arXiv:2106.04554 (2021).

[336] J. Xu, D. Moyer, P. E. Grant, P. Golland, J. E. Iglesias, E. Adalsteinsson, SVoRT: iterative transformer for slice-to-volume registration in fetal brain MRI, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2022, pp. 3–13.

[337] Y. Rong, M. Rosu-Bubulac, S. H. Benedict, Y. Cui, R. Ruo, T. Connell, R. Kashani, K. Latifi, Q. Chen, H. Geng, et al., Rigid and deformable image registration for radiation therapy: a self-study evaluation guide for NRG oncology clinical trial participation, Practical Radiation Oncology 11 (4) (2021) 282–298.

[338] J. Chen, Y. He, E. Frey, Y. Li, Y. Du, ViT-V-Net: Vision transformer for unsupervised volumetric medical image registration, in: Medical Imaging with Deep Learning, 2021.
URL https://openreview.net/forum?id=h3HC1EU7AEz

[339] J. Chen, E. C. Frey, Y. He, W. P. Segars, Y. Li, Y. Du, Transmorph: Transformer for unsupervised medical image registration, Medical image analysis 82 (2022) 102615.

[340] Y. Zhang, Y. Pei, H. Zha, Learning dual transformer network for diffeomorphic registration, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 129–138.

[341] J. Shi, Y. He, Y. Kong, J.-L. Coatrieux, H. Shu, G. Yang, S. Li, XMorpher: Full transformer for deformable medical image registration via cross attention, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2022, pp. 217–226.

[342] K. Miao, A. Gokul, R. Singh, S. Petryk, J. Gonzalez, K. Keutzer, T. Darrell, Prior knowledge-guided attention in self-supervised vision transformers, arXiv preprint arXiv:2209.03745 (2022).

[343] T. C. Mok, A. Chung, Affine medical image registration with coarse-to-fine vision transformer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 20835–20844.

[344] F. Milletari, N. Navab, S.-A. Ahmadi, V-net: Fully convolutional neural networks for volumetric medical image segmentation, in: 2016 fourth international conference on 3D vision (3DV), IEEE, 2016, pp. 565–571.

[345] M. Jaderberg, K. Simonyan, A. Zisserman, et al., Spatial transformer networks, Advances in neural information processing systems 28 (2015).

[346] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, R. 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) (2007) 1498–1507.

[347] W. Segars, J. Bond, J. Frush, S. Hon, C. Eckersley, C. H. Williams, J. Feng, D. J. Tward, J. Ratnanather, M. Miller, et al., Population of anatomically variable 4D XCAT adult phantoms for imaging research and optimization, Medical physics 40 (4) (2013) 043701.

[348] X. Zhuang, J. Shen, Multi-scale patch and multi-modality atlases for whole heart segmentation of MRI, Medical image analysis 31 (2016) 77–87.

[349] R. Gharleghi, D. G. Samarasinghe, P. A. Sowmya, D. S. Beier, Automated segmentation of coronary arteries (Mar. 2020). doi:10.5281/zenodo.3819799. URL https://doi.org/10.5281/zenodo.3819799

[350] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K.-i. Komatsu, M. Matsui, H. Fujita, Y. Kodera, K. 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) (2000) 71–74.

[351] D. W. Shattuck, M. Mirza, V. Adisetiyo, C. Hojatkashani, G. Salamon, K. L. Narr, R. A. Poldrack, R. M. Bilder, A. W. Toga, Construction of a 3D probabilistic atlas of human cortical structures, Neuroimage 39 (3) (2008) 1064–1080.

[352] K. Payette, P. de Dumast, H. Kebiri, I. Ezhov, J. C. Paetzold, S. Shit, A. Iqbal, R. Khan, R. Kottke, P. Grehten, et al., An automatic multi-tissue human fetal brain segmentation benchmark using the fetal tissue annotation dataset, Scientific Data 8 (1) (2021) 1–14.

[353] B. D. De Vos, F. F. Berendsen, M. A. Viergever, H. Sokooti, M. Staring, I. Išgum, A deep learning framework for unsupervised affine and deformable image registration, Medical image analysis 52 (2019) 128–143.

[354] S. Zhao, T. Lau, J. Luo, I. Eric, C. Chang, Y. Xu, Unsupervised 3D end-to-end medical image registration with volume tweening network, IEEE journal of biomedical and health informatics 24 (5) (2019) 1394–1404.

[355] Y. Xiong, B. Du, P. Yan, Reinforced transformer for medical image captioning, in: International Workshop on Machine Learning in Medical Imaging, Springer, 2019, pp. 673–680.

[356] J. Zhang, Y. Nie, J. Chang, J. J. Zhang, Surgical instruction generation with transformers, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 290–299.

[357] Z. Chen, Y. Song, T.-H. Chang, X. Wan, Generating radiology reports via memory-driven transformer, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 1439–1449.

[358] D. You, F. Liu, S. Ge, X. Xie, J. Zhang, X. 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, Springer, 2021, pp. 72–82.

[359] F. Nooralahzadeh, N. P. Gonzalez, T. Frauenfelder, K. Fujimoto, M. Krauthammer, Progressive transformer-based generation of radiology reports, arXiv preprint arXiv:2102.09777 (2021).

[360] A. Yan, Z. He, X. Lu, J. Du, E. Chang, A. Gentili, J. McAuley, C.-n. Hsu, Weakly supervised contrastive learning for chest X-Ray report generation, in: Findings of the Association for Computational Linguistics: EMNLP 2021, 2021, pp. 4009–4015.

[361] Z. Chen, Y. Shen, Y. Song, X. Wan, Cross-modal memory networks for radiology report generation, arXiv preprint arXiv:2204.13258 (2022).

[362] C. Y. Li, X. Liang, Z. Hu, E. P. Xing, Knowledge-driven encode, retrieve, paraphrase for medical image report generation, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33, 2019, pp. 6666–6673.

[363] F. Liu, X. Wu, S. Ge, W. Fan, Y. 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, 2021, pp. 13753–13762.

[364] J. Lovelace, B. Mortazavi, Learning to generate clinically coherent chest x-ray reports, in: Findings of the Association for Computational Linguistics: EMNLP 2020, 2020, pp. 1235–1243.

[365] G. Liu, Y. Liao, F. Wang, B. Zhang, L. Zhang, X. Liang, X. Wan, S. Li, Z. Li, S. Zhang, et al., Medical-vlbert: Medical visual language bert for covid-19 ct report generation with alternate learning, IEEE Transactions on Neural Networks and Learning Systems 32 (9) (2021) 3786–3797.

[366] O. Alfarghaly, R. Khaled, A. Elkorany, M. Helal, A. Fahmy, Automated radiology report generation using conditioned transformers, Informatics in Medicine Unlocked 24 (2021) 100557.

[367] H. T. Nguyen, D. Nie, T. Badamdorj, Y. Liu, Y. Zhu, J. Truong, L. Cheng, Automated generation of accurate & fluent medical X-ray reports, arXiv preprint arXiv:2108.12126 (2021).

[368] Y. Wang, Z. Lin, J. Tian, Z. Shi, Y. Zhang, J. Fan, Z. He, Confidence-guided radiology report generation, arXiv preprint arXiv:2106.10887 (2021).

[369] M. Endo, R. Krishnan, V. Krishna, A. Y. Ng, P. Rajpurkar, Retrieval-based chest X-ray report generation using a pre-trained contrastive language-image model, in: Machine Learning for Health, PMLR, 2021, pp. 209–219.

[370] S. Bannur, S. Hyland, Q. Liu, F. Perez-Garcia, M. Ilse, D. C. Castro, B. Boecking, H. Sharma, K. Bouzid, A. Thieme, et al., Learning to exploit temporal structure for biomedical vision-language processing, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2023, pp. 15016–15027.

[371] S. Wang, Z. Zhao, X. Ouyang, Q. Wang, D. Shen, Chatcad: Interactive computer-aided diagnosis on medical image using large language models, arXiv preprint arXiv:2302.07257 (2023).

[372] B. Jing, P. Xie, E. Xing, On the automatic generation of medical imaging reports, arXiv preprint arXiv:1711.08195 (2017).

[373] M. Stefanini, M. Cornia, L. Baraldi, S. Cascianelli, G. Fiameni, R. Cucchiara, From show to tell: a survey on deep learning-based image captioning, IEEE Transactions on Pattern Analysis and Machine Intelligence (2022).

[374] D. Demner-Fushman, M. D. Kohli, M. B. Rosenman, S. E. Shooshan, L. Rodriguez, S. Antani, G. R. Thoma, C. J. McDonald, Preparing a collection of radiology examinations for distribution and retrieval, Journal of the American Medical Informatics Association 23 (2) (2016) 304–310.

[375] M. M. A. Monshi, J. Poon, V. Chung, Deep learning in generating radiology reports: A survey, Artificial Intelligence in Medicine 106 (2020) 101878.

[376] R. Vedantam, C. Lawrence Zitnick, D. Parikh, Cider: Consensus-based image description evaluation, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 4566–4575.

[377] Y. Li, X. Liang, Z. Hu, E. P. Xing, Hybrid retrieval-generation reinforced agent for medical image report generation, Advances in neural information processing systems 31 (2018).

[378] Y. Zhang, X. Wang, Z. Xu, Q. Yu, A. Yuille, D. Xu, When radiology report generation meets knowledge graph, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34, 2020, pp. 12910–12917.

[379] G. Huang, Z. Liu, L. Van Der Maaten, K. Q. Weinberger, Densely connected convolutional networks, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 4700–4708.

[380] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, Bleu: a method for automatic evaluation of machine translation, in: Proceedings of the 40th annual meeting of the Association for Computational Linguistics, 2002, pp. 311–318.

[381] E. Rojas-Muñoz, K. Couperus, J. Wachs, Daisi: Database for ai surgical instruction, arXiv preprint arXiv:2004.02809 (2020).

[382] S. Banerjee, A. Lavie, METEOR: An automatic metric for mt evaluation with improved correlation with human judgments, in: Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization, 2005, pp. 65–72.

[383] C.-Y. Lin, Rouge: A package for automatic evaluation of summaries, in: Text summarization branches out, 2004, pp. 74–81.

[384] P. Anderson, B. Fernando, M. Johnson, S. Gould, Spice: Semantic propositional image caption evaluation, in: European conference on computer vision, Springer, 2016, pp. 382–398.

[385] A. E. Johnson, T. J. Pollard, N. R. Greenbaum, M. P. Lungren, C.-y. Deng, Y. Peng, Z. Lu, R. G. Mark, S. J. Berkowitz, S. Horng, MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs, arXiv preprint arXiv:1901.07042 (2019).

[386] J. Ni, C.-N. Hsu, A. Gentili, J. McAuley, Learning visual-semantic embeddings for reporting abnormal findings on chest X-rays, arXiv preprint arXiv:2010.02467 (2020).

[387] G. Liu, Y. Liao, Z. Li, Covid-19ct dataset, https://covid19ct.github.io/.

[388] M. Li, R. Liu, F. Wang, X. Chang, X. Liang, Auxiliary signal-guided knowledge encoder-decoder for medical report generation, World Wide Web (2022) 1–18.

[389] A. Smit, S. Jain, P. Rajpurkar, A. Pareek, A. Y. Ng, M. P. Lungren, CheXbert: combining automatic labelers and expert annotations for accurate radiology report labeling using BERT, arXiv preprint arXiv:2004.09167 (2020).

[390] P. Messina, P. Pino, D. Parra, A. Soto, C. Besa, S. Uribe, M. Andía, C. Tejos, C. Prieto, D. Capurro, A survey on deep learning and explainability for automatic report generation from medical images, ACM Computing Surveys (CSUR) 54 (10s) (2022) 1–40.

[391] F. Liu, X. Ren, G. Zhao, X. Sun, Layer-wise cross-view decoding for sequence-to-sequence learning, arXiv preprint arXiv:2005.08081 (2020).

[392] F. Liu, S. Ge, X. Wu, Competence-based multimodal curriculum learning for medical report generation, arXiv preprint arXiv:2206.14579 (2022).

[393] R. H. Zhang, Q. Liu, A. X. Fan, H. Ji, D. Zeng, F. Cheng, D. Kawahara, S. Kurohashi, Minimize exposure bias of seq2seq models in joint entity and relation extraction, arXiv preprint arXiv:2009.07503 (2020).

[394] S. J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, V. Goel, Self-critical sequence training for image captioning, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2017, pp. 7008–7024.

[395] X. Wang, Y. Chen, W. Zhu, A survey on curriculum learning, IEEE Transactions on Pattern Analysis and Machine Intelligence (2021).

[396] M. Cornia, M. Stefanini, L. Baraldi, R. Cucchiara, Meshed-memory transformer for image captioning, in: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020, pp. 10578–10587.

[397] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, L. Zettlemoyer, Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension, arXiv preprint arXiv:1910.13461 (2019).

[398] P. Rajpurkar, J. Irvin, K. Zhu, B. Yang, H. Mehta, T. Duan, D. Ding, A. Bagul, C. Langlotz, K. Shpanskaya, et al., Chexnet: Radiologist-level pneumonia detection on chest X-rays with deep learning, arXiv preprint arXiv:1711.05225 (2017).

[399] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al., Language models are unsupervised multitask learners, OpenAI blog 1 (8) (2019) 9.

[400] W. Su, X. Zhu, Y. Cao, B. Li, L. Lu, F. Wei, J. Dai, VL-BERT: Pre-training of generic visual-linguistic representations, in: International Conference on Learning Representations, 2020.
URL https://openreview.net/forum?id=SygXPaEYvH

[401] V. Ramesh, N. A. Chi, P. Rajpurkar, Improving radiology report generation systems by removing hallucinated references to non-existent priors, in: Machine Learning for Health, PMLR, 2022, pp. 456–473.

[402] J. Zhou, X. He, L. Sun, J. Xu, X. Chen, Y. Chu, L. Zhou, X. Liao, B. Zhang, X. Gao, SkinGPT-4: An interactive dermatology diagnostic system with visual large language model, medRxiv (2023) 2023–06.

[403] C. Ma, Z. Wu, J. Wang, S. Xu, Y. Wei, Z. Liu, L. Guo, X. Cai, S. Zhang, T. Zhang, et al., ImpressionGPT: an iterative optimizing framework for radiology report summarization with chatGPT, arXiv preprint arXiv:2304.08448 (2023).

[404] OpenAI, Chatgpt: Optimizing language models for dialogue (2023).
URL https://openai.com/blog/chatgpt/

[405] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Advances in neural information processing systems 33 (2020) 1877–1901.

[406] T. Lin, Y. Wang, X. Liu, X. Qiu, A survey of transformers, AI Open (2022).

[407] J. Pavlopoulos, V. Kougia, I. Androutsopoulos, D. Papamichail, Diagnostic captioning: a survey, Knowledge and Information Systems (2022) 1–32.

[408] X. Meng, C. H. Ganoe, R. T. Sieberg, Y. Y. Cheung, S. Hassanpour, Self-supervised contextual language representation of radiology reports to improve the identification of communication urgency, AMIA Summits on Translational Science Proceedings 2020 (2020) 413.

[409] G. Alicioglu, B. Sun, A survey of visual analytics for explainable artificial intelligence methods, Computers & Graphics 102 (2022) 502–520.

[410] A. Singh, S. Sengupta, V. Lakshminarayanan, Explainable deep learning models in medical image analysis, Journal of Imaging 6 (6) (2020) 52.

[411] B. Hou, G. Kaissis, R. M. Summers, B. Kainz, Ratchet: Medical transformer for chest X-ray diagnosis and reporting, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 293–303.

[412] A. Binder, G. Montavon, S. Lapuschkin, K.-R. Müller, W. Samek, Layer-wise relevance propagation for neural networks with local renormalization layers, in: International Conference on Artificial Neural Networks, Springer, 2016, pp. 63–71.

[413] S. Kim, J. Nam, B. C. Ko, Vit-net: Interpretable vision transformers with neural tree decoder, in: International Conference on Machine Learning, PMLR, 2022, pp. 11162–11172.

[414] Z. Li, F. Liu, W. Yang, S. Peng, J. Zhou, A survey of convolutional neural networks: analysis, applications, and prospects, IEEE transactions on neural networks and learning systems (2021).

[415] J. Gu, H. Kwon, D. Wang, W. Ye, M. Li, Y.-H. Chen, L. Lai, V. Chandra, D. Z. Pan, Multi-scale high-resolution vision transformer for semantic segmentation,

in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 12094–12103.

[416] Y. Lee, J. Kim, J. Willette, S. J. Hwang, MPViT: Multipath vision transformer for dense prediction, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 7287–7296.

[417] H. Reynaud, A. Vlontzos, B. Hou, A. Beqiri, P. Leeson, B. Kainz, Ultrasound video transformers for cardiac ejection fraction estimation, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 495–505.

[418] Y. Long, Z. Li, C. H. Yee, C. F. Ng, R. H. Taylor, M. Unberath, Q. 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, Springer, 2021, pp. 415–425.

[419] T. Czempiel, M. Paschali, D. Ostler, S. T. Kim, B. Busam, N. Navab, Opera: Attention-regularized transformers for surgical phase recognition, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 604–614.

[420] Z. Zhao, Y. Jin, P.-A. Heng, TraSeTR: track-to-segment transformer with contrastive query for instance-level instrument segmentation in robotic surgery, in: 2022 International Conference on Robotics and Automation (ICRA), IEEE, 2022, pp. 11186–11193.

[421] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, H. Hu, Video swin transformer, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 3202–3211.

[422] M. Ding, B. Xiao, N. Codella, P. Luo, J. Wang, L. Yuan, Davit: Dual attention vision transformers, in: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV, Springer, 2022, pp. 74–92.

[423] A. Ali, H. Touvron, M. Caron, P. Bojanowski, M. Douze, A. Joulin, I. Laptev, N. Neverova, G. Synnaeve, J. Verbeek, et al., Xcit: Cross-covariance image transformers, Advances in neural information processing systems 34 (2021) 20014–20027.

[424] M. Maaz, A. Shaker, H. Cholakkal, S. Khan, S. W. Zamir, R. M. Anwer, F. Shahbaz Khan, Edgenext: efficiently amalgamated cnn-transformer architecture for mobile vision applications, in: European Conference on Computer Vision, Springer, 2022, pp. 3–20.

[425] Z. Shen, M. Zhang, H. Zhao, S. Yi, H. Li, Efficient attention: Attention with linear complexities, in: Proceedings of the IEEE/CVF winter conference on applications of computer vision, 2021, pp. 3531–3539.

[426] Z. Zhang, U. Bagci, Dynamic linear transformer for 3D biomedical image segmentation, in: Machine Learning in Medical Imaging: 13th International Workshop, MLMI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 18, 2022, Proceedings, Springer, 2022, pp. 171–180.

[427] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, K. H. Maier-Hein, nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation, Nature methods 18 (2) (2021) 203–211.

[428] C. Gros, A. Lemay, O. Vincent, L. Rouhier, A. Bucquet, J. P. Cohen, J. Cohen-Adad, Ivadomed: A medical imaging deep learning toolbox, arXiv preprint arXiv:2010.09984 (2020).