

Deep Lesion Tracker: Monitoring Lesions in 4D Longitudinal Imaging Studies

Jinzheng Cai¹ Youbao Tang¹ Ke Yan¹ Adam P. Harrison¹ Jing Xiao² Gigin Lin³ Le Lu¹
¹PAII Inc. ²Ping An Technology ³Chang Gung Memorial Hospital

Abstract

Monitoring treatment response in longitudinal studies plays an important role in clinical practice. Accurately identifying lesions across serial imaging follow-up is the core to the monitoring procedure. Typically this incorporates both image and anatomical considerations. However, matching lesions manually is labor-intensive and time-consuming. In this work, we present deep lesion tracker (DLT), a deep learning approach that uses both appearance- and anatomical-based signals. To incorporate anatomical constraints, we propose an anatomical signal encoder, which prevents lesions being matched with visually similar but spurious regions. In addition, we present a new formulation for Siamese networks that avoids the heavy computational loads of 3D cross-correlation. To present our network with greater varieties of images, we also propose a self-supervised learning (SSL) strategy to train trackers with unpaired images, overcoming barriers to data collection. To train and evaluate our tracker, we introduce and release the first lesion tracking benchmark, consisting of 3891 lesion pairs from the public DeepLesion database. The proposed method, DLT, locates lesion centers with a mean error distance of 7mm. This is 5% better than a leading registration algorithm while running 14 times faster on whole CT volumes. We demonstrate even greater improvements over detector or similarity-learning alternatives. DLT also generalizes well on an external clinical test set of 100 longitudinal studies, achieving 88% accuracy. Finally, we plug DLT into an automatic tumor monitoring workflow where it leads to an accuracy of 85% in assessing lesion treatment responses, which is only 0.46% lower than the accuracy of manual inputs.

1. Introduction

Monitoring treatment response by identifying and measuring corresponding lesions is critical in radiological workflows [20, 40, 1, 47]. Manually conducting these procedures is labor-intensive, as expert clinicians must review multiple images and go back and forth between these images for comparison. This is usually subject to considerable inter-

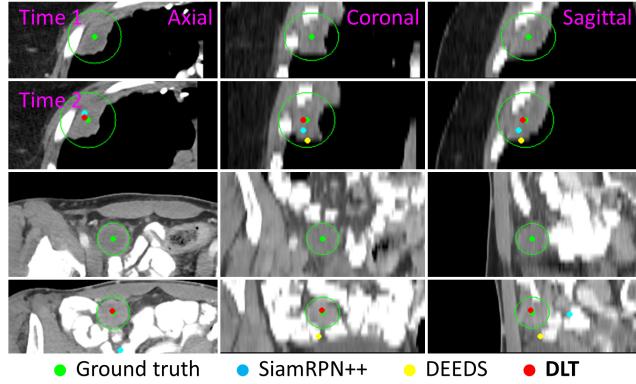


Figure 1. Comparison of our approach with two state-of-the-art approaches for 3D tracking. The proposed DLT can predict lesion centers more precisely than SiamRPN++ [29] and DEEDS [22].

observer variability [48]. Therefore, computer aided tools have the opportunity to lower costs, increase turnaround speeds, and improve reliability.

Automatic image-based lesion monitoring can be decomposed into several sub-procedures: (1) detect lesions of interest; (2) then track instances of the same lesion across different time points; and (3) measure changes among the identified instances. The first step of detecting lesions of interest can be formulated as object detection. In general, the computer vision field has made progress toward this problem [21, 31, 65]. However, medical imaging has its distinct challenges as the data is often in 3D format, e.g., computed tomography (CT), and usually the required annotations are unavailable. Therefore, there are efforts to improve object detection with medical images [27, 45, 55, 12, 58, 49, 44]. Similarly step (3) also has many viable solutions because it can be formulated as (3D) object segmentation, which is a fundamental topic that attracts attentions from both computer vision [34, 56, 15] and medical image analysis [42, 38, 17, 43, 11, 50]. In contrast, step (2), tracking the same lesion across different time points, is not as well developed as lesion detection and segmentation. Part of the lack of development can be attributed to the lack of good benchmark datasets to evaluate performance. In this work, we address this by both introducing a public benchmark and also formulating a powerful lesion tracking solution, called DLT, that can accurately match instances of the same lesion

across different images captured at different time points and contrast phases by using both appearance and anatomical signals. In Fig. 1, we show two real-life examples of lesion tracking.

Similar with visual tracking in the general computer vision, lesion tracking can be viewed as to match instances of the same lesion in neighboring time frames. However, it is challenging due to changes in size and appearance. Lesion size can enlarge multiple times than its baseline or nadir. Meanwhile, its appearance varies during the follow-up exam because of morphological or functional changes, commonly attributed to necrosis or changes in vascularity. Therefore, an effective tracker should handle both size and visual changes of lesions. Trackers based on image registration [1, 47] are robust to appearance changes, as registration inherently introduces anatomical constraints for lesion matching. The involved body part and surrounding organs of the target lesion are constrained among different images. However, registration algorithms [22, 23, 36, 37, 5] are usually less sensitive to local image changes; thus, they can be inaccurate to track small-sized lesions or lesions with large shape changes. On the other hand, appearance-based trackers [40, 14] handle size and appearance changes by projecting lesion images into an embedding space [62, 59], where images of the same lesion have similar embeddings and images of different lesions are different from one another. However, these appearance-based trackers may mismatch lesions with visually similar but spurious backgrounds. Therefore, to combine the merits of both strategies, we design our tracker to conduct appearance based recognition under anatomical constraints.

Because the proposed deep lesion tracker (DLT) is a deep learning model, providing enough training data is a prerequisite for good performance. To this end, we construct a dataset with 3891 lesion pairs, collected from DeepLesion [61], to train and evaluate different tracking solutions. We publicly release this dataset to facilitate related research¹. Although more training pairs can promote a stronger tracker, labor and time costs preclude easily collecting and annotating a large number of longitudinal studies for a specific clinical application. Therefore, we also introduce an effective self-supervised learning (SSL) strategy to train trackers. Importantly, this strategy can train lesion trackers using images from only one time point, meaning non-longitudinal datasets can be used, which are more readily collected. This allows for a more ready introduction of more lesion instances with varied appearances and sizes.

With the proposed DLT and model training strategies, we achieve 89% matching accuracy on a test set of 480 lesion pairs. Meanwhile, we demonstrate that DLT is robust to inaccurate tracking initializations, *i.e.*, the given initial lesion center. In our robustness study, inaccurate initializa-

tion causes 10% accuracy drops on SiamRPN++ [29] and DEEDS [22]. In contrast, the accuracy of DLT only decreases by 1.9%. We then apply DLT to an external testing set of 100 real-life clinical longitudinal studies, delivering 88% matching accuracy and demonstrating excellent generalizability. Finally, we plug DLT into a lesion monitoring pipeline to simulate automatic treatment monitoring. The workflow assesses lesion treatment responses with 85% accuracy, which is only 0.46% lower than the accuracy of manual inputs.

2. Related Work

Visual object tracking is an active research topic in general computer vision [8, 35, 18, 54, 53, 6, 39, 51, 52]. We focus our review on recent progresses, especially deep learning based approaches.

Tracking as Similarity Learning. Tracking of target objects can be achieved via similarity comparisons between the object template and proposals from the search domain. Similarities are measured by either color/intensity representations [19], spatial configurations [63, 33], or their combinations [6]. Recently, deep learning features are more widely used for visual tracking [53, 39, 18, 54] as they outperform hand-crafted features with more expressive representations. To efficiently extract and compare deep learning features, SiamFC [7] and CFNet [52] use a cross-correlation layer at the end of Siamese architectures [9]. This cross-correlation layer uses Siamese feature maps extracted from the template image patch as a kernel to operate fully circular convolution on the corresponding Siamese feature maps of the search image. This procedure encodes the information regarding the relative position of the target object inside the search image. Within the same framework of SiamFC, SiamRPN++ [29] introduced strategies to allow training of Siamese networks with modern very deep networks, *e.g.*, dense convolutional network (DenseNet) [25], to further boost tracking accuracy. This is critical for medical image analysis as many medical applications lack large-scale training data and rely on transfer learning of pre-trained networks for good performance [46].

Siamese networks have also been investigated in medical image analysis. Gomariz *et al.* [14] applied 2D Siamese networks to track liver landmarks in ultra-sound videos. Liu *et al.* [32] extended similar 2D Siamese networks in a coarse-to-fine fashion. While, Rafael-Palou *et al.* [40] performed 3D Siamese networks with CT series, only shallow network architectures were evaluated on tracking lung nodules. However, we follow SiamRPN++ [29] to use Siamese networks with 3D DenseNet backbones and apply it to conduct universal lesion tracking in whole body CT images. Processing different types of lesions with a unified deep learning model [45, 49, 61, 62, 59, 60, 10, 12] demonstrates computational efficiency and could also alleviate model over-

¹<https://github.com/JimmyCai91/DLT>

fitting. Different from prior formulations of Siamese networks, we propose a simple but effective 3D kernel decomposition to speed up 3D cross-correlation operations for object matching. This provides dramatic boosts in efficiency, reducing over 65% of FLOPs in our fast cross-correlation (FCC) layer.

Tracking as Detector Learning. Tracking as detector learning relies on developing discriminative models to separate the target from background regions [2, 3, 24, 19]. A discriminative model that is suitable for visual tracking should consist of two core components, namely a classifier that can be efficiently updated online during visual tracking [2, 3, 24] and a powerful feature representation, *e.g.* features extracted by convolutional neural networks (CNNs) [28, 25] that can let the classifier easily differentiate objects in the feature space. Following this strategy, SO-DLT [54], FCNT [53], and MDNet [39] all train CNNs offline from large-scale object recognition tasks so that the learnt feature representation is general with visual objects. During tracking, they freeze the lower layers of the network as a feature extractor and update the higher layers to adapt to the specific video domain.

In this work, we consider the strategy of tracking via detector learning and accordingly construct our strong lesion tracking baselines. Given the specialty of processing medical data, especially 4D CT images (3D image plus time), there are no baseline methods ready for comparison. Thus, we construct our own lesion tracking baselines by concatenating the state-of-the-art lesion detection [12, 58] models with deep learning feature extractors [62, 59]. However, the tracker developed with this strategy can be sub-optimal since the detection models feature extractors are developed from independent offline tasks. In contrast, our proposed DLT unifies the tasks of feature extraction and target object localization in an end-to-end structure and outperforms these detector learning baselines with higher accuracy and faster speed.

Tracking Priors from Image Registration. Visual tracking in video follows a prior of spatial consistency, which means the search space in the next video frame can be constrained to be near to the current location. This prior is helpful for improving tracking efficiency and making the model robust to background distractors [51, 7, 14]. Similarly, lesion tracking in CT should follow a spatial consistency governed by anatomical considerations. This implies that the surrounding organs and structures of a lesion will not drastically change. Under such constraints, image registration approaches [22, 23, 36, 37, 5] can perform lesion tracking via image alignment. Specifically, registration algorithms are designed to optimize the global structural alignment, *i.e.* accurately align boundaries of large organs, while being robust to local changes. Nonetheless, although reported results suggest that registration algorithms are use-

ful for aligning large-sized lesions [47, 41, 64], they can fail to track small-sized lesions and struggle whenever there are local changes in the lesion’s appearance.

In this work, we improve upon the capabilities of registration approaches using deep learning based lesion appearance recognition to match lesions based on both visual and anatomical signals. Specifically, we first roughly initialize the location of a target lesion using image registration, *i.e.*, affine registration [36]. Then, our DLT deep learning model refines the location to the lesion center using appearance-based cues. In contrast with approaches that use the spatial and structural priors simply in pre- [51, 7] or post-processing [14], DLT takes them as its inputs and propagates them together with CT-based visual signal to generate the final target location. The priors also function as attention guidance, letting the appearance learning focus on vital image regions.

3. Deep Lesion Tracker

We build DLT based on the structure of Siamese networks because they are efficient and deliver state-of-the-art visual tracking performance for many computer vision tasks. The core component of Siamese-based tracking is a correlation filter, which is also known as cross-correlation layer. It uses Siamese features extracted from the template image patch as a kernel to perform explicit convolutional scanning over the entire extent of the search image feature maps. Fig. 2 shows its overall configuration. Our goal is to apply the proposed model to process three dimensional medical data, *i.e.*, CT images. Therefore, we create network backbones in 3D and introduce an anatomy signal encoder (ASE) to guide lesion tracking with anatomical constraints. To avoid the prohibitive computational expenses of 3D cross-correlation between the template and the search image, we introduce a simple but effective formulation to speed up this procedure.

Problem definition. We use I_t and I_s to respectively denote a template and search CT image. In I_t , a lesion is known with its center μ_t and radius r_t . Given I_t , I_s , μ_t , and r_t , the task of lesion tracking is to locate the same lesion in I_s by predicting its new center μ_s .

3.1. Image Encoder: 3D DenseFPN

In lesion tracking, the Siamese network needs to process lesions with varied appearances and sizes in 3D images. As shown in Fig. 3, we use a deep 3D image encoder with large model capacity, so that it can learn effective feature representations. Specifically, we transform DenseNet into 3D by duplicating its 2D convolutional kernels along the third direction and then downscaling weight values by the number of duplications [13]. This configuration is found to be more effective than 3D UNet [17] on modeling universal lesion appearances [12]. We then add a feature pyramid network

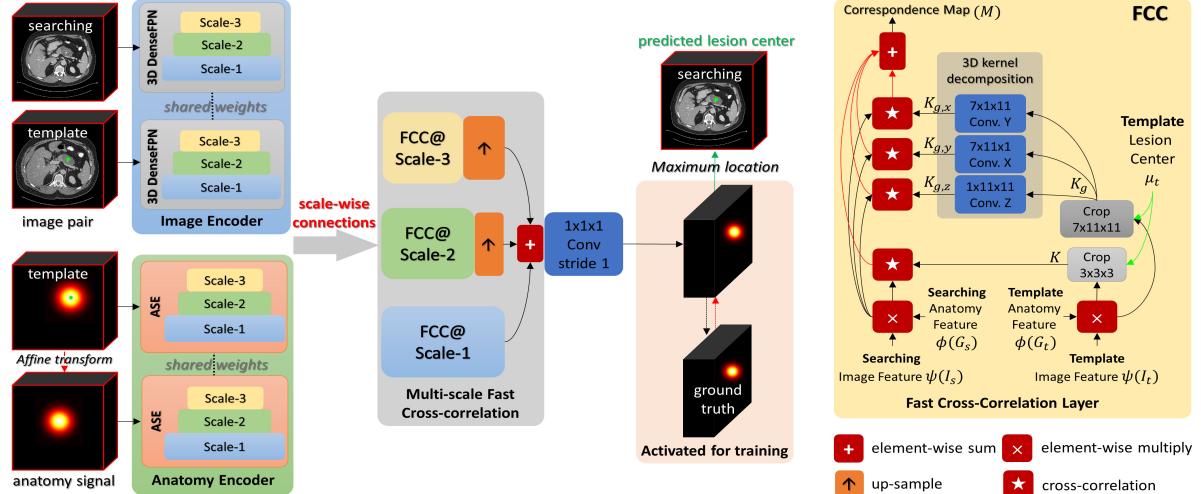


Figure 2. The configuration of our proposed deep lesion tracker.

(FPN) [30] after the 3D DenseNet to generate visual features at three scales. We visually depict the detailed configuration of 3D DenseFPN in Fig. 3. For clarity, we use ψ_1 , ψ_2 , and ψ_3 to refer to the image mapping functions that generate feature maps from the largest to the smallest resolutions, respectively.

3.2. Anatomy Signal Encoder and Its Inputs

We observe that directly implementing lesion tracking with Siamese networks can produce matches with visually similar but spurious regions. In contrast, affine registration [36] is a robust approach to roughly align CT images. It is achieved by solving

$$\mathcal{T}_{\text{Aff}} = \arg \min_{\mathcal{T}_{\text{Aff}} \in \mathcal{A}} \|\mathcal{T}_{\text{Aff}}(I_t) - I_s\|_1, \quad (1)$$

where \mathcal{A} is the space of affine transforms. The projected location of the template lesion, $\mathcal{T}_{\text{Aff}}(\mu_t)$, is usually located close to the actual target lesion. While prior art has used affine registration as pre- [51, 7] or post-processing [14], these do not provide mechanisms for incorporation into a tracking pipeline that cross-correlates template features across the entire extent of the search image. For example, pre-registering will have minimal effect on the translation-invariant cross-correlation. Instead, as shown in Fig. 2, we encode anatomy signals as Gaussian heatmaps centered at lesion locations:

$$G(\mu, nr) = \exp \left(-\frac{\sum_{i \in \{x, y, z\}} (i - \mu^i)^2}{2(nr)^2} \right), \quad (2)$$

where we find $n = 4$ delivers the best performance. For I_t we simply use the template lesion location and size: $G(\mu_t, nr_t)$. For I_s we use the affine-projected location and size of the template lesion: $G(\mathcal{T}_{\text{Aff}}(\mu_t), n\mathcal{T}_{\text{Aff}}(r_t))$. For clarity, we simply refer to the template and search anatomy sig-

nal maps as G_t and G_s , respectively. We solve Eq. 1 using SimpleElastix [36].

Fig. 3 depicts the network configuration of the proposed ASE. It encodes anatomical signals into high-dimensional anatomical features with three different resolutions. In correspondence with 3D DenseFPN, we denote the network functions for the three scales as ϕ_1 , ϕ_2 , and ϕ_3 from the largest to the smallest, respectively.

3.3. Fast Cross-Correlation

As mentioned, correlation is a core operation of Siamese-based tracking, which creates a correspondence map between target and search features, $\psi(I_t)$ and $\phi(G_t)$, respectively. Because we perform the same operation at each scale, we drop the scale subscripts here for simplicity. To conduct cross-correlation, we first fuse image and anatomy features. For example, to fuse $\psi(I_t)$ and $\phi(G_t)$ we use

$$F = \psi(I_t) \odot \phi(G_t), \quad (3)$$

where \odot is element-wise multiplication and we constrain $\phi(G_t)$ to have the same shape as $\psi(I_t)$. We observe from experiments that fusing $\psi(I_t)$ and $\phi(G_t)$ with \odot performs better than channel-wise concatenation. Next, we define a cropping function to extract a $3 \times 3 \times 3$ template kernel as,

$$K = \mathcal{C}(F, \mu_t, (3, 3, 3)), \quad (4)$$

where the kernel is centered at μ_t after any potential feature downscaling. To encode the global image context better, we also extract another larger size kernel $K_g = \mathcal{C}(F, \mu_t, (7, 11, 11))$. Here we limit its size in the z -direction to be 7 since the size of I_t during model training is only $(32, 384, 384)$.

Following the traditional cross-correlation operation [7], we define the correspondence map as,

$$M = (K \star S) + (K_g \star S), \quad (5)$$

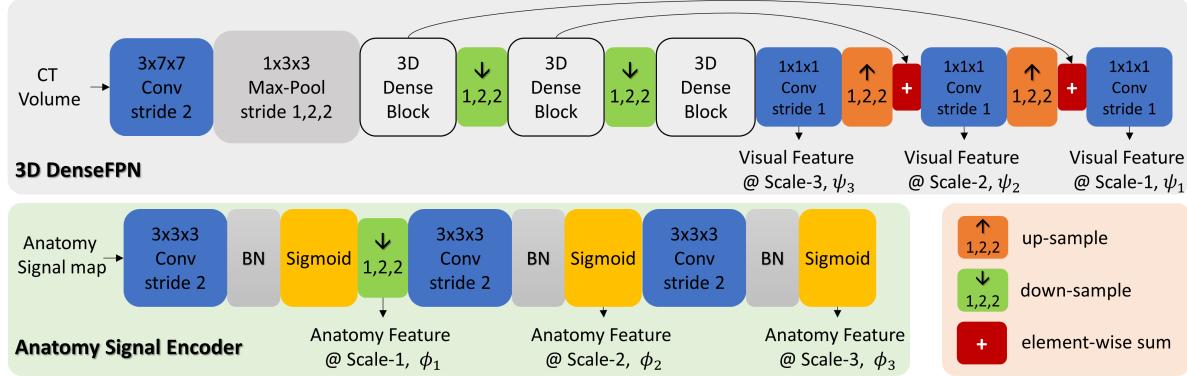


Figure 3. Network configurations of the proposed image encoder 3D DenseFPN and anatomy signal encoder (ASE).

where $S = \psi(I_s) \odot \phi(G_s)$ and $+$ is the element-wise sum. Unfortunately, a direct use of K_g introduces a heavy computational load. We propose to decompose K_g along the axial, coronal, and sagittal directions and obtain flattened kernels as $K_{g,z} \in \mathbf{R}^{(1,11,11)}$, $K_{g,x} \in \mathbf{R}^{(7,1,11)}$, and $K_{g,y} \in \mathbf{R}^{(7,11,1)}$, where we omit the dimensions of batch size for clarity. As Fig. 2 demonstrates, the proposed FCC layer performs the flattening using learned 3D convolutions configured to produce an output of identical size as the kernel, except with one dimension flattened. The resulting faster version of Eq. 5 is

$$M = (K \star S) + \sum_{i \in x,y,z} K_{g,i} \star S. \quad (6)$$

We also tested kernel decomposition by simply extracting the middle ‘‘slices’’ of K_g along the three dimensions, but it did not perform as well as the learned flattening operations.

Adding back the scale subscripts, the final output is a probability map:

$$\hat{Y} = \sigma(W^T(M_1 + U_2 + U_3) + b), \quad (7)$$

where $\sigma(\cdot)$ is the Sigmoid function, W and b are parameters of the final fully convolutional layer, U_2 is M_2 up-scaled by $(1, 2, 2)$, and U_3 is M_3 up-scaled by $(1, 4, 4)$. The predicted lesion center μ_p is the index of the global maximum in \hat{Y} .

4. Supervised and Self-Supervised Learning

DLT is capable of both supervised and self-supervised learning (SSL). It is flexible enough to learn from paired annotations, when enough are available, and to also use efficient self-supervised learning.

4.1. Supervised Learning

Based on the introduced network architecture, \hat{Y} , the output of DLT is a dense probability map representing the likelihood of each location to be the target lesion center. Therefore, we define the ground truth as a Gaussian kernel centered at the target location μ_s . Formally, we first define

$Y = \mathcal{G}(\mu_s, r_s)$ and then downsize it to match the dimensions of \hat{Y} . We use focal loss [31, 65] in training:

$$\mathcal{L}_{sl} = \sum_i \begin{cases} (1 - \hat{y}_i)^\alpha \log(\hat{y}_i) & \text{if } y_i = 1 \\ (1 - y_i)^\beta (\hat{y}_i)^\alpha \log(1 - \hat{y}_i) & \text{otherwise} \end{cases}, \quad (8)$$

where y_i and \hat{y}_i are the i -th voxels in Y and \hat{Y} , respectively, and $\alpha = 2$ and $\beta = 4$ are focal-loss hyper-parameters [31, 65]. The ground-truth heat map is < 1 everywhere except at the lesion center voxel. So that the training can converge quickly, it ignores hard voxels that are near μ_s .

Center augmentation. In practice, labels from clinicians may not represent the exact lesion centers. The provided location, μ_t , may shift inside the central area. Therefore, to increase model robustness we train DLT with random location shifts. This is achieved by adding μ_t with $\Delta\mu_t$, which is randomly sampled from the sphere $\|\Delta\mu_t\|_2 \leq 0.25r_t$.

4.2. Self-Supervised Learning

Since our proposed DLT is built upon Siamese pairwise comparison, it inherently supports learning with self-supervision. The key insight is that effective visual representation for object recognition can be learned by comparing the template image, I_t , with its augmented counterparts. With I_t , we implement data augmentations including (1) elastic deformations at random scales ranging from 0 to 0.25, (2) rotations in the xy -plane with a random angle ranging from -10 to 10 degrees, (3) random scales ranging from 0.75 to 1.25, (4) random crops, (5) add Gaussian noise with zero mean and a random variance ranging from 0 to 0.05, and (6) Gaussian blurring with a random sigma ranging from 0.5 to 1.5 [26]. Each augmentation individually takes place with the probability of 0.5. For clarity, we define \mathcal{T}_{aug} as any combination of the data augmentations. Therefore, each self-supervised image ‘‘pair’’ comprises I_t and $\mathcal{T}_{aug}(I_t)$ with corresponding anatomical signals of G_t and $\mathcal{T}_{aug}(G_t)$. The same training procedure as supervised learning can then be followed. It is worth mentioning that our

SSL strategy shares a similar spirit with recent contrastive learning studies that matches an image with its transformed version [16], but in the pixel-level.

We select non-longitudinal images from DeepLesion [61] and use the bounding box annotations as μ_t and r_t . When bounding box annotations are not available, the template lesions can be extracted by applying a pre-trained universal lesion detector on I_t and randomly selecting top-scoring proposals. However, we do not explore that here.

Limited by GPU memory, when combining the supervised learning with SSL, we switch the training of DLT between both schemes as:

$$\mathcal{L}_{mix} = \begin{cases} \mathcal{L}_{ssl} & \text{if } \lambda \leq \tau \\ \mathcal{L}_{sl} & \text{otherwise} \end{cases}, \quad (9)$$

where $\lambda \in [0, 1]$ is a random number and we empirically set the threshold τ to 0.25 in our experiments.

5. Experiments

5.1. Datasets

DeepLesion is a large-scale CT database of lesions released by the National Institute of Health (NIH) in 2018 [61, 62]. It contains over 30 thousand lesions and each lesion is associated with a size measurement defined by the response evaluation criteria in solid tumours (RECIST) [20]. The RECIST measurement consists of two diameters: the longest diameter followed by the longest diameter that is perpendicular to the first one. Both diameters are drawn by doctor in a manually selected axial slice. Based on this measurement, we define the ground truth lesion center μ to be the mean of diameters' four end points and the radius r is approximated to be the half of the longest diameters. In total, our publicly released deep longitudinal study (DLS) dataset inherits 3008, 403, and 480 lesion pairs from the DeepLesion's train, validate, and test splits, respectively.

From Chang Gung Memorial Hospital (IRB 202000584A3C601), we also collected an external validation set that consists of 536 lesions from 100 longitudinal studies of 86 patients, including 45 cervical, 27 endometrial, and 14 ovarian cancers (mean age, 53.3 years). The median time interval between CT studies was 217 days (range, 8-2304 days). We apply the best DLT configuration, developed on the DeepLesion dataset, to track the corresponding target lesions, if they exist. To assess the tracking accuracy we measured the acceptance rate of a board-certified radiologist with over 10 years of clinical practice experience.

5.2. Evaluation Metrics

For an annotated pair of lesion a and b , we evaluate tracking both from a to b and from b to a . Therefore, in

total, we have 906 and 960 directed lesion pairs in the validation and test sets, respectively. We define a center point matching (CPM) accuracy, which represents the percentage of correctly matched lesions. A match will be counted correct when the Euclidean distance between the ground truth center and the predicted center is smaller than a threshold. We first set the threshold to be the corresponding lesion radius and refer the matching accuracy CPM@Radius or simply CPM. However this threshold is not tight enough to differentiate trackers as some lesions have large sizes. We then use an adaptive threshold $\min(r, 10mm)$ to limit the allowed maximum offset in large lesions and we refer to this matching accuracy as CPM@10mm. We empirically use 10mm because 55% lesions in the test set have larger than 10mm radius.

We also measure the absolute offset between ground truth and predicted centers in mm and report the mean Euclidean distance (MED) and its projections MED_X, MED_Y, MED_Z in each direction. The speed of trackers is counted using seconds per volume (spv).

5.3. Comparisons with State-of-the-art Approaches

Traditional registration approaches. We use both the widely used rigid affine registration method [36] and DEEDS [22] deformable registration. The latter is considered the state-of-the-art deformable approach for CT registration [57]. The implementation that we use is optimized in C++ [23] and the CT volumes have been resampled to the isotropic resolution of 2mm.

Learning based registration. We use VoxelMorph [5], which is a general deep learning framework for deformable medical image registration that can deliver state-of-the-art performance with a much faster speeds than traditional approaches. We train VoxelMorph with image pairs from DLS. Image pairs are first aligned by affine registration and then resampled to 0.8mm by 0.8mm in xy -plane with a slice thickness of 2mm. The same image resolution is applied to all of the following experiments.

Tracking by detector learning. These approaches first detect lesion candidates. Then, an image encoder is used to project both the template lesion and the detected candidates into feature vectors. Lastly, a nearest neighbor classifier is applied to identify the tracked lesion. We tested the detector with the 2D LENS [61] and 3D VULD [12] detectors, both of which report good performance on DeepLesion. As for the image encoder, we tested LesionGraph [62] and LesaNet [59], which are also developed from DeepLesion for lesion attribute description. Therefore, we evaluate four baselines, *i.e.*, LENS-LesionGraph, LENS-LesaNet, VULD-LesionGraph, and VULD-LesaNet.

Tracking by similarity learning. We adapt SiamRPN++ [29] with 3D DenseFPN so that it can process CT images and perform fair comparison with DLT. The largest size of

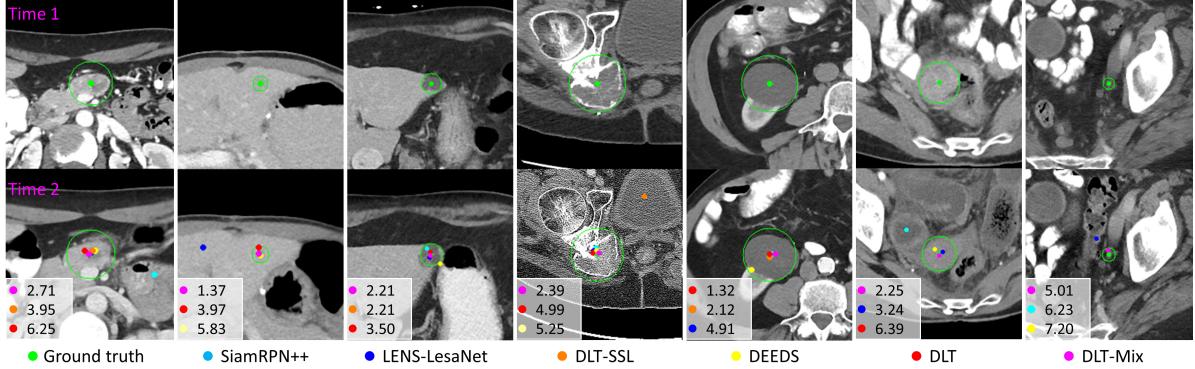


Figure 4. Comparisons of our methods with three state-of-the-art trackers. The top 3 closest to center distances are reported in mm.

Method	CPM@ 10mm	CPM@ Radius	MED _X (mm)	MED _Y (mm)	MED _Z (mm)	MED (mm)	speed (spv)
Affine [36]	48.33	65.21	4.1±5.0	5.4±5.6	7.1±8.3	11.2±9.9	1.82
VoxelMorph [4]	49.90	65.59	4.6±6.7	5.2±7.9	6.6±6.2	10.9±10.9	0.46
LENS-LesionGraph [58, 62]	63.85	80.42	2.6±4.6	2.7±4.5	6.0±8.6	8.0±10.1	4.68
VULD-LesionGraph [12, 62]	64.69	76.56	3.5±5.2	4.1±5.8	6.1±8.8	9.3±10.9	9.07
VULD-LesaNet [12, 59]	65.00	77.81	3.5±5.3	4.0±5.7	6.0±8.7	9.1±10.8	9.05
SiamRPN++ [29]	68.85	80.31	3.8±4.8	3.8±4.8	4.8±7.5	8.3±9.2	2.24
LENS-LesaNet [58, 59]	70.00	84.58	2.7±4.8	2.6±4.7	5.7±8.6	7.8±10.3	4.66
DLT-SSL	71.04	81.52	3.8±5.3	3.7±5.5	5.4±8.4	8.8±10.5	3.57
DEEDS [22]	71.88	85.52	2.8±3.7	3.1±4.1	5.0±6.8	7.4±8.1	15.3
DLT-Mix	78.65	88.75	3.1±4.4	3.1±4.5	4.2±7.6	7.1±9.2	3.54
DLT	78.85	86.88	3.5±5.6	2.9±4.9	4.0±6.1	7.0±8.9	3.58

Table 1. Comparisons between the proposed DLT and state-of-the-art approaches.

Method	CPM@10mm	MED (mm)
SiamRPN++ [29]	51.27 (\downarrow 17.6)	10.6±10.3 (\uparrow 2.3)
DEEDS [22]	53.85 (\downarrow 18.0)	9.8±8.9 (\uparrow 2.4)
DLT-SSL	64.24 (\downarrow 6.80)	10.0±11.4 (\uparrow 1.2)
DLT	70.36 (\downarrow 8.49)	8.1±8.7 (\uparrow 1.2)
DLT-Mix	75.03 (\downarrow 3.62)	8.0±10.5 (\uparrow 0.9)

Table 2. Robustness evaluation. \downarrow and \uparrow demonstrate decrease and increase of measurements, respectively, compared with the values reported in Table 1.

the template kernel is (3, 5, 5) for computational efficiency. **DLT and its variants.** DLT is trained using DLS. DLT-SSL is trained using only SSL with non-longitudinal training images of DeepLesion that do not exist in DLS. DLT-Mix is trained with a combination of supervised and self-supervised learning, which is defined by Eq. 9.

Results. Table 1 shows the comparative results. With CPM@Radius, DLT and DLT-Mix achieve the first and second places, respectively, leading DEEDS at the third place by over 6%. DLT-SSL is at the 4th place outperforming its SSL counterparts, *i.e.*, affine registration and VoxelMorph, by over 20%. With CPM@Radius, DLT-Mix is the best tracker, and it outperforms DEEDS and SiamRPN++ by 3.2% and 8.4%, respectively. With MED, DLT performs the best. We notice that LENS-LesionGraph outperforms DLT in MED_X by 0.9mm because LENS is a 2D lesion detector with a bounding-box regression layer, which is dedicated to locating the lesion accurately in the

xy-plane. Similarly, LENS-LesaNet outperforms DLT by 0.3mm in MED_Y. However, in MED_Z, DLT greatly outperforms LENS-LesionGraph and LENS-LesaNet by 2mm and 1.7mm, respectively, showing the importance of 3D DenseFPN. In terms of speed, affine registration and VoxelMorph are the top 2 methods but they are not as accurate as the others. Among the top 3 methods, DLT and DLT-Mix run about 4 times faster than DEEDS on the DeepLesion dataset. Fig. 4 shows seven visual examples of lesion tracking, where the results produced by our trackers are closer to the ground truth than others.

Robustness evaluation. In this experiment, we simulate human inputs. In testing, we shift the template center μ_t with $\Delta\mu_t$, which is randomly sampled from the sphere $\|\Delta\mu\|_2 \leq 0.25r_t$. For each directed lesion pair, 9 shifted centers together with the original center are stored. In total, we create 9060 and 9600 directed lesion pairs from the validation and test sets, respectively. With these augmented lesion pairs, we evaluate trackers to see if they are robust with inaccurate human inputs or not.

Table 2 shows the results. DLT-Mix is in the first place for both CPM and MED metrics. DEEDS turns out to be the most vulnerable method with 18% drop in CPM and 2.4mm increase in MED. In comparison, DLT-Mix only drops 3.62% in CPM and increases only 0.9mm in MED. Additionally, DLT-SSL is more robust than DLT in CPM, demonstrating the benefit of SSL in training robust trackers.

id	Eq. 6: K_g		ψ, ϕ dim.	Eq. 2: G size (n)	test MED	speed spv
	size	learn				
a	N/A	N/A	64	4	9.3	1.44
b	7,7,7	✓	64	4	9.4	2.38
c	7,15,15	✓	64	4	7.7	24.1
d	7,11,11	✓	64	2	7.4	3.51
e	7,11,11	✓	64	8	8.5	3.51
f	7,11,11	✓	32	4	8.7	2.25
g	7,11,11	✓	128	4	7.9	5.83
h	7,11,11	✓	64	N/A	9.3	3.51
i	7,11,11	✗	64	4	9.3	3.51
j	7,11,11	✓	64	4	7.9	3.51

Table 3. Parameter analysis of the proposed components.

Parameter Analysis. Table 3 presents our parameter analysis for different model configurations, with model j representing our final configuration without the multiplication fusion of Eq. 3 or the center augmentation of Sec. 4.1. We present test results, but note that our model selection was based off of our validation (which can be found in the supplementary material). Model a is identical to our final model, except that the global kernel has been disabled, resulting in significant MED increases and demonstrating the importance of the global kernel. Models b and c explore different global kernel sizes, indicating performance can vary somewhat, but is not overly sensitive to the choice. However, too large of a kernel results in an order of magnitude greater runtime, justifying our choice of a (7, 11, 11) kernel. As model e demonstrates, when the ASE heat map of Eq. 2 covers too large of an area it can lose its specificity, resulting in performance degradation. Models f and g show the effect of different embedding feature dimensions, again showing that performance is not overly sensitive to this choice, as long as the embedding dimension is large enough. In terms of the need for the anatomy signal of ASE, model h demonstrates its removal considerably increases the MED. Finally, model i 's performance shows that the learnable decomposition of Eq. 6 is critical for accurate tracking. Adding Eq. 3 and center augmentation to model j results in our final configuration featured in Table 1.

5.4. Impact on Downstream Measurements

In this experiment, we compare trackers with downstream size measurements. We use a pre-trained model, OneClick [50] that takes the image I_s and the predicted lesion center μ_p as its inputs and regresses the RECIST diameters of the target lesion. For simplicity, we only compare the long diameters. We use evaluation metrics including mean absolute error (MAE) in mm, growth accuracy, and treatment response accuracy. With the template diameter d_t , search diameter d_s , and OneClick predicted diameter d_p , we define d_p as a correct growth prediction, if and only if the inequality $(d_s - d_t)(d_p - d_t) > 0$ holds. The growth accuracy represents the percentage of correct growth predictions. The treatment response, $\rho = (d_s - d_t)/d_t$, is defined based on the

Input generator	MAE (mm)	Growth acc. (%)	Response acc. (%)
DEEDS [22]	2.69±4.12	78.02	84.17
DLT	2.47±3.58	79.69	85.10
Manual inputs	2.31±3.16	79.69	85.56

Table 4. Impact on automatic lesion size measurement when using the OneClick [50] model.

Method	CPM@Radius	speed (spv)
DEEDS [22]	85.6	67.1±17.8
DLT	88.4	4.7±0.35

Table 5. External evaluation.

RECIST guideline [20], which classifies a treatment response as partial response if $\rho \leq -0.3$, as progressive disease if $\rho \geq 0.2$, or as stable disease if $\rho \in (-0.3, 0.2)$. We then predict treatment response using $\rho_p = (d_p - d_t)/d_t$.

We tested DLT, DEEDS, and manual inputs, *i.e.* the ground truth lesion centers. Table 4 shows the results. DLT outperforms DEEDS in MAE by 0.22mm, which is an 8% improvement. Compared with manual inputs, DLT exhibits the same growth accuracy and is only 0.46% lower in the treatment response accuracy.

External Evaluation. We further invite a board-certified radiologist to manually assess DLT with 100 longitudinal studies recruited from real-life clinical workflows. The user provides binarized responses, *i.e.*, *inside-* or *outside-lesion* for the CPM@Radius metric. We compared the tracking results of DLT with DEEDS in Table 5. DLT delivers 88.4% CPM accuracy and outperforms DEEDS by 2.8%. Besides, DLT requires only 4.67 seconds to process a whole body CT, which is over 14 times faster than DEEDS. These results also underscore the value of our DLS dataset.

6. Conclusion & Discussion

In this work, we introduce a new public benchmark for lesion tracking and present DLT as our solution. Due to the different setup of medical applications, DLT differs from general visual trackers in two aspects. First, DLT does not regress bounding boxes for target lesions because as mentioned in Sec. 5.4, the lesion size can be accurately predicted by the down stream measurement module. Second, DLT does not perform long-term tracking because time points in longitudinal studies is much less than general videos. Also, manual calibration occurs much more often in lesion tracking than general object tracking.

Our presented DLT has been demonstrated effective for lesion tracking, outperforming a comprehensive set of baselines that represent various tracking strategies. DLT can be trained via either supervised or self-supervised learning, where the combination of both training schemes results in the best performance and robustness. We benchmark the task of lesion tracking on our DLS dataset which will be made available upon request.

References

- [1] Diego Ardila, Atilla P Kiraly, Sujeeth Bharadwaj, Bokyung Choi, Joshua J Reicher, Lily Peng, Daniel Tse, Mozziyar Etemadi, Wenxing Ye, Greg Corrado, et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nature medicine*, 25(6):954–961, 2019. [1](#), [2](#)
- [2] Shai Avidan. Support vector tracking. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(8):1064–1072, 2004. [3](#)
- [3] Boris Babenko, Ming-Hsuan Yang, and Serge J. Belongie. Robust object tracking with online multiple instance learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(8):1619–1632, 2011. [3](#)
- [4] Guha Balakrishnan, Amy Zhao, Mert Sabuncu, John Guttag, and Adrian V. Dalca. An unsupervised learning model for deformable medical image registration. *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9252–9260, 2018. [7](#)
- [5] Guha Balakrishnan, Amy Zhao, Mert R. Sabuncu, John V. Guttag, and Adrian V. Dalca. Voxelmorph: A learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging*, 38(8):1788–1800, 2019. [2](#), [3](#), [6](#)
- [6] Luca Bertinetto, Jack Valmadre, Stuart Golodetz, Ondrej Miksik, and Philip H. S. Torr. Staple: Complementary learners for real-time tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1401–1409. 2016. [2](#)
- [7] Luca Bertinetto, Jack Valmadre, João F. Henriques, Andrea Vedaldi, and Philip H. S. Torr. Fully-convolutional siamese networks for object tracking. In *Eur. Conf. Comput. Vis. Worksh.*, pages 850–865, 2016. [2](#), [3](#), [4](#)
- [8] David S. Bolme, J. Ross Beveridge, Bruce A. Draper, and Yui Man Lui. Visual object tracking using adaptive correlation filters. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2544–2550. 2010. [2](#)
- [9] Jane Bromley, James W. Bentz, Léon Bottou, Isabelle Guyon, Yann LeCun, Cliff Moore, Eduard Säckinger, and Roopak Shah. Signature verification using A "siamese" time delay neural network. *Int. J. Pattern Recognit. Artif. Intell.*, 7(4):669–688, 1993. [2](#)
- [10] J. Cai, A. P. Harrison, Y. Zheng, K. Yan, Y. Huo, J. Xiao, L. Yang, and L. Lu. Lesion-harvester: Iteratively mining unlabeled lesions and hard-negative examples at scale. *IEEE Trans. Med. Imaging*, pages 1–1, 2020. [2](#)
- [11] Jinzheng Cai, Youbao Tang, Le Lu, Adam P. Harrison, Ke Yan, Jing Xiao, Lin Yang, and Ronald M. Summers. Accurate weakly-supervised deep lesion segmentation using large-scale clinical annotations: Slice-propagated 3d mask generation from 2d RECIST. In *Medical Image Computing and Computer Assisted Intervention*, pages 396–404. 2018. [1](#)
- [12] Jinzheng Cai, Ke Yan, Chi-Tung Cheng, Jing Xiao, Chien-Hung Liao, Le Lu, and Adam P. Harrison. Deep volumetric universal lesion detection using light-weight pseudo 3d convolution and surface point regression. In *Medical Image Computing and Computer Assisted Intervention*, pages 3–13. 2020. [1](#), [2](#), [3](#), [6](#), [7](#)
- [13] João Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4724–4733. 2017. [3](#)
- [14] Alvaro Gomariz Carrillo, Weiye Li, Ece Ozkan, Christine Tanner, and Orcun Goksel. Siamese networks with location prior for landmark tracking in liver ultrasound sequences. In *IEEE Int. Symposium on Biomedical Imaging*, pages 1757–1760. 2019. [2](#), [3](#), [4](#)
- [15] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(4):834–848, 2018. [1](#)
- [16] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *CoRR*, abs/2002.05709, 2020. [6](#)
- [17] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In *Medical Image Computing and Computer Assisted Intervention*, pages 424–432, 2016. [1](#), [3](#)
- [18] Martin Danelljan, Gustav Häger, Fahad Shahbaz Khan, and Michael Felsberg. Convolutional features for correlation filter based visual tracking. In *Int. Conf. Comput. Vis. Worksh.*, pages 621–629. 2015. [2](#)
- [19] Martin Danelljan, Fahad Shahbaz Khan, Michael Felsberg, and Joost van de Weijer. Adaptive color attributes for real-time visual tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1090–1097. 2014. [2](#), [3](#)
- [20] E. Eisenhauer, P. Therasse, J. Bogaerts, and et al. New response evaluation criteria in solid tumours: revised recist guideline (version 1.1). *European journal of cancer*, 45(2):228–247, 2009. [1](#), [6](#), [8](#)
- [21] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 580–587. 2014. [1](#)
- [22] M. P. Heinrich, M. Jenkinson, M. Brady, and J. A. Schnabel. Mrf-based deformable registration and ventilation estimation of lung ct. *IEEE Trans. Med. Imaging*, 32(7):1239–1248, 2013. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [23] Matthias P. Heinrich, Oskar Maier, and Heinz Handels. Multi-modal multi-atlas segmentation using discrete optimisation and self-similarities. In *IEEE Int. Symposium on Biomedical Imaging*, pages 27–30. 2015. [2](#), [3](#), [6](#)
- [24] João F. Henriques, Rui Caseiro, Pedro Martins, and Jorge P. Batista. Exploiting the circulant structure of tracking-by-detection with kernels. In *Eur. Conf. Comput. Vis.*, pages 702–715. 2012. [3](#)
- [25] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2261–2269. 2017. [2](#), [3](#)
- [26] Fabian Isensee, Jens Petersen, André Klein, David Zimmerer, Paul F. Jaeger, Simon Kohl, Jakob Wasserthal, Gregor Koehler, Tobias Norajitra, Sebastian J. Wirkert, and Klaus H. Maier-Hein. nnunet: Self-adapting framework for u-net-

- based medical image segmentation. *CoRR*, abs/1809.10486, 2018. 5
- [27] Chenhan Jiang, Shaoju Wang, Xiaodan Liang, Hang Xu, and Nong Xiao. Elixirnet: Relation-aware network architecture adaptation for medical lesion detection. In *AAAI*, pages 11093–11100. 2020. 1
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Adv. Neural Inform. Process. Syst.*, pages 1106–1114, 2012. 3
- [29] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 4282–4291. 2019. 1, 2, 6, 7
- [30] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 936–944. 2017. 4
- [31] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(2):318–327, 2020. 1, 5
- [32] Fei Liu, Dan Liu, Jie Tian, Xiaoyan Xie, Xin Yang, and Kun Wang. Cascaded one-shot deformable convolutional neural networks: Developing a deep learning model for respiratory motion estimation in ultrasound sequences. *Med. Image Anal.*, 65:101793, 2020. 2
- [33] Ting Liu, Gang Wang, and Qingxiong Yang. Real-time part-based visual tracking via adaptive correlation filters. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4902–4912. 2015. 2
- [34] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3431–3440. 2015. 1
- [35] Chao Ma, Jia-Bin Huang, Xiaokang Yang, and Ming-Hsuan Yang. Hierarchical convolutional features for visual tracking. In *Int. Conf. Comput. Vis.*, pages 3074–3082. 2015. 2
- [36] Kasper Marstal, Floris F. Berendsen, Marius Staring, and Stefan Klein. Simpleelastix: A user-friendly, multi-lingual library for medical image registration. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 574–582. 2016. 2, 3, 4, 6, 7
- [37] Shun Miao, Z. Jane Wang, and Rui Liao. A CNN regression approach for real-time 2d/3d registration. *IEEE Trans. Med. Imaging*, 35(5):1352–1363, 2016. 2, 3
- [38] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *3DV*, pages 565–571. 2016. 1
- [39] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4293–4302. 2016. 2, 3
- [40] Xavier Rafael-Palou, Anton Aubanell, Ilaria Bonavita, Mario Ceresa, Gemma Piella, Vicent Ribas, and Miguel A. González Ballester. Re-identification and growth detection of pulmonary nodules without image registration using 3d siamese neural networks. *Med. Image Anal.*, 67:101823, 2021. 1, 2
- [41] Ashwin Raju, Chi-Tung Cheng, Yuankai Huo, Jinzheng Cai, Junzhou Huang, Jing Xiao, Le Lu, Chien-Hung Liao, and Adam P. Harrison. Co-heterogeneous and adaptive segmentation from multi-source and multi-phase CT imaging data: A study on pathological liver and lesion segmentation. In *Eur. Conf. Comput. Vis.*, pages 448–465. 2020. 3
- [42] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer Assisted Intervention*, pages 234–241. 2015. 1
- [43] Holger R. Roth, Le Lu, Amal Farag, Hoo-Chang Shin, Jiamin Liu, Evrim B. Turkbey, and Ronald M. Summers. Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In *Medical Image Computing and Computer Assisted Intervention*, pages 556–564. 2015. 1
- [44] Holger R. Roth, Le Lu, Ari Seff, Kevin M. Cherry, Joanne Hoffman, Shijun Wang, Jiamin Liu, Evrim Turkbey, and Ronald M. Summers. A new 2.5d representation for lymph node detection using random sets of deep convolutional neural network observations. In *Medical Image Computing and Computer Assisted Intervention*, pages 520–527. 2014. 1
- [45] Qingbin Shao, Lijun Gong, Kai Ma, Hualuo Liu, and Yefeng Zheng. Attentive CT lesion detection using deep pyramid inference with multi-scale booster. In *Medical Image Computing and Computer Assisted Intervention*, pages 301–309. 2019. 1, 2
- [46] Hoo-Chang Shin, Holger R. Roth, Mingchen Gao, Le Lu, Ziyue Xu, Isabella Nogues, Jianhua Yao, Daniel J. Mollura, and Ronald M. Summers. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans. Med. Imaging*, 35(5):1285–1298, 2016. 2
- [47] Maxine Tan, Zheng Li, Yuchen Qiu, Scott D. McMeekin, Theresa C. Thai, Kai Ding, Kathleen N. Moore, Hong Liu, and Bin Zheng. A new approach to evaluate drug treatment response of ovarian cancer patients based on deformable image registration. *IEEE Trans. Med. Imaging*, 35(1):316–325, 2016. 1, 2, 3
- [48] Youbao Tang, Adam P Harrison, Mohammadhdadi Bagheri, Jing Xiao, and Ronald M Summers. Semi-automatic RECIST labeling on CT scans with cascaded convolutional neural networks. In *Medical Image Computing and Computer Assisted Intervention*, pages 405–413. 2018. 1
- [49] Youbao Tang, Ke Yan, Yuxing Tang, Jiamin Liu, Jin Xiao, and Ronald M. Summers. Uldor: A universal lesion detector for ct scans with pseudo masks and hard negative example mining. In *IEEE Int. Symposium on Biomedical Imaging*, pages 833–836. 2019. 1, 2
- [50] Youbao Tang, Ke Yan, Jing Xiao, and Ronald M. Summers. One click lesion RECIST measurement and segmentation on CT scans. In *Medical Image Computing and Computer Assisted Intervention*, pages 573–583. 2020. 1, 8

- [51] Ran Tao, Efstratios Gavves, and Arnold W. M. Smeulders. Siamese instance search for tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1420–1429. 2016. 2, 3, 4
- [52] Jack Valmadre, Luca Bertinetto, João F. Henriques, Andrea Vedaldi, and Philip H. S. Torr. End-to-end representation learning for correlation filter based tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5000–5008. 2017. 2
- [53] Lijun Wang, Wanli Ouyang, Xiaogang Wang, and Huchuan Lu. Visual tracking with fully convolutional networks. In *Int. Conf. Comput. Vis.*, pages 3119–3127. 2015. 2, 3
- [54] Naiyan Wang, Siyi Li, Abhinav Gupta, and Dit-Yan Yeung. Transferring rich feature hierarchies for robust visual tracking. *CoRR*, abs/1501.04587, 2015. 2, 3
- [55] Xudong Wang, Zhaowei Cai, Dashan Gao, and Nuno Vasconcelos. Towards universal object detection by domain attention. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7289–7298. 2019. 1
- [56] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. *Int. J. Comput. Vis.*, 125(1-3):3–18, 2017. 1
- [57] Zhoubing Xu, Christopher P. Lee, Matthias P. Heinrich, Marc Modat, Daniel Rueckert, Sébastien Ourselin, Richard G. Abramson, and Bennett A. Landman. Evaluation of six registration methods for the human abdomen on clinically acquired CT. *IEEE Trans. Biomed. Eng.*, 63(8):1563–1572, 2016. 6
- [58] Ke Yan, Jinzheng Cai, Youjing Zheng, Adam P. Harrison, Dakai Jin, Youvao Tang, Yuxing Tang, Lingyun Huang, Jing Xiao, and Le Lu. Learning from multiple datasets with heterogeneous and partial labels for universal lesion detection in CT. *CoRR*, abs/2009.02577, 2020. 1, 3, 7
- [59] Ke Yan, Yifan Peng, Veit Sandfort, Mohammad Bagheri, Zhiyong Lu, and Ronald M. Summers. Holistic and comprehensive annotation of clinically significant findings on diverse CT images: Learning from radiology reports and label ontology. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8523–8532. 2019. 2, 3, 6, 7
- [60] Ke Yan, Youbao Tang, Yifan Peng, Veit Sandfort, Mohammad Bagheri, Zhiyong Lu, and Ronald M. Summers. MULAN: multitask universal lesion analysis network for joint lesion detection, tagging, and segmentation. In *Medical Image Computing and Computer Assisted Intervention*, pages 194–202. 2019. 2
- [61] Ke Yan, Xiaosong Wang, Le Lu, and Ronald M. Summers. Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *J. Med. Imaging*, 5(3), 2018. 2, 6
- [62] Ke Yan, Xiaosong Wang, Le Lu, Ling Zhang, Adam P. Harrison, Mohammad Bagheri, and Ronald M. Summers. Deep lesion graphs in the wild: Relationship learning and organization of significant radiology image findings in a diverse large-scale lesion database. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9261–9270. 2018. 2, 3, 6, 7
- [63] Rui Yao, Qinfeng Shi, Chunhua Shen, Yanning Zhang, and Anton van den Hengel. Part-based visual tracking with online latent structural learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2363–2370. 2013. 2
- [64] Ling Zhang, Yu Shi, Jiawen Yao, Yun Bian, Kai Cao, Dakai Jin, Jing Xiao, and Le Lu. Robust pancreatic ductal adenocarcinoma segmentation with multi-institutional multi-phase partially-annotated CT scans. In *Medical Image Computing and Computer Assisted Intervention*, pages 491–500. 2020. 3
- [65] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *CoRR*, abs/1904.07850, 2019. 1, 5