

Supplementary Material

In this material, we provide some additional illustrations of the paper. Sec. 1 visualizes the DeepLesion dataset and describes some details. Sec. 2 provides implementation details of the self-supervised body-part regressor. More content-based lesion retrieval results are presented in Sec. 3. Sec. 4 illustrates the intra-patient lesion matching task and the intra-patient lesion graph.

1. DeepLesion Dataset: Visualization and Details

To provide an overview of the DeepLesion dataset, we draw a scatter map to show the distribution of the types and relative body locations of the lesions in Fig. 3. From the lesion types and sample images, one can see that the relative body locations of the lesions are consistent with their actual physical positions, proving the validity of the location information used in the paper, particularly the self-supervised body-part regressor. Some lesion types like bone

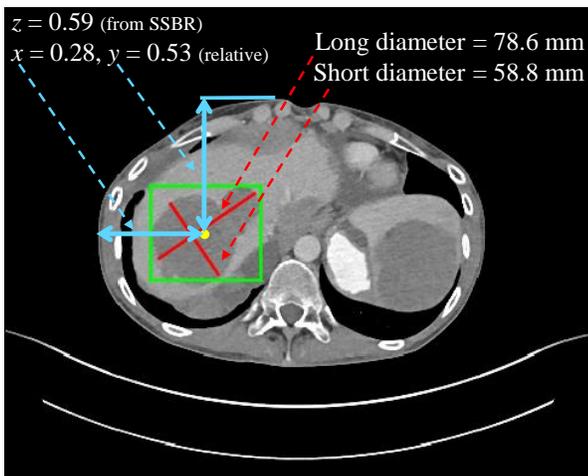


Figure 1. Location and size of a sample lesion. The red lines are the long and short diameters annotated by radiologists during their daily work. The green box is the bounding box calculated from the diameters. The yellow dot is the center of the bounding box. The blue lines indicate the relative x - and y -coordinates of the lesion. The z -coordinate is predicted by SSBR. Best viewed in color.

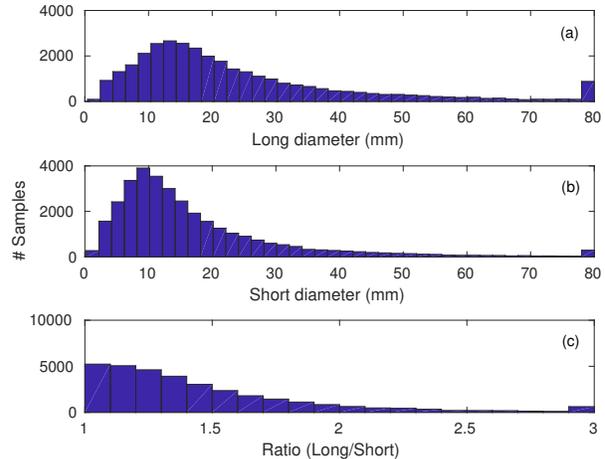


Figure 2. Distribution of the lesion-sizes in DeepLesion. For clarity, values greater than the upper bound of the x -axis of each plot are grouped in the last bin of each histogram.

and soft tissue have widespread locations. Neighboring types such as lung/mediastinum and abdomen/liver/kidney have large overlap in location due to inter-subject variabilities. Besides, we can clearly see the considerable diversity of DeepLesion.

Fig. 1 illustrates the approach to obtain the location and size of a lesion. In order to locate a lesion in the body, we first obtain the mask of the body in the axial slice, then compute the relative position (0–1) of the lesion center to get the x - and y -coordinates. As for z , the self-supervised body-part regressor (SSBR) is used. We also show the distribution of the lesion-sizes in Fig. 2.

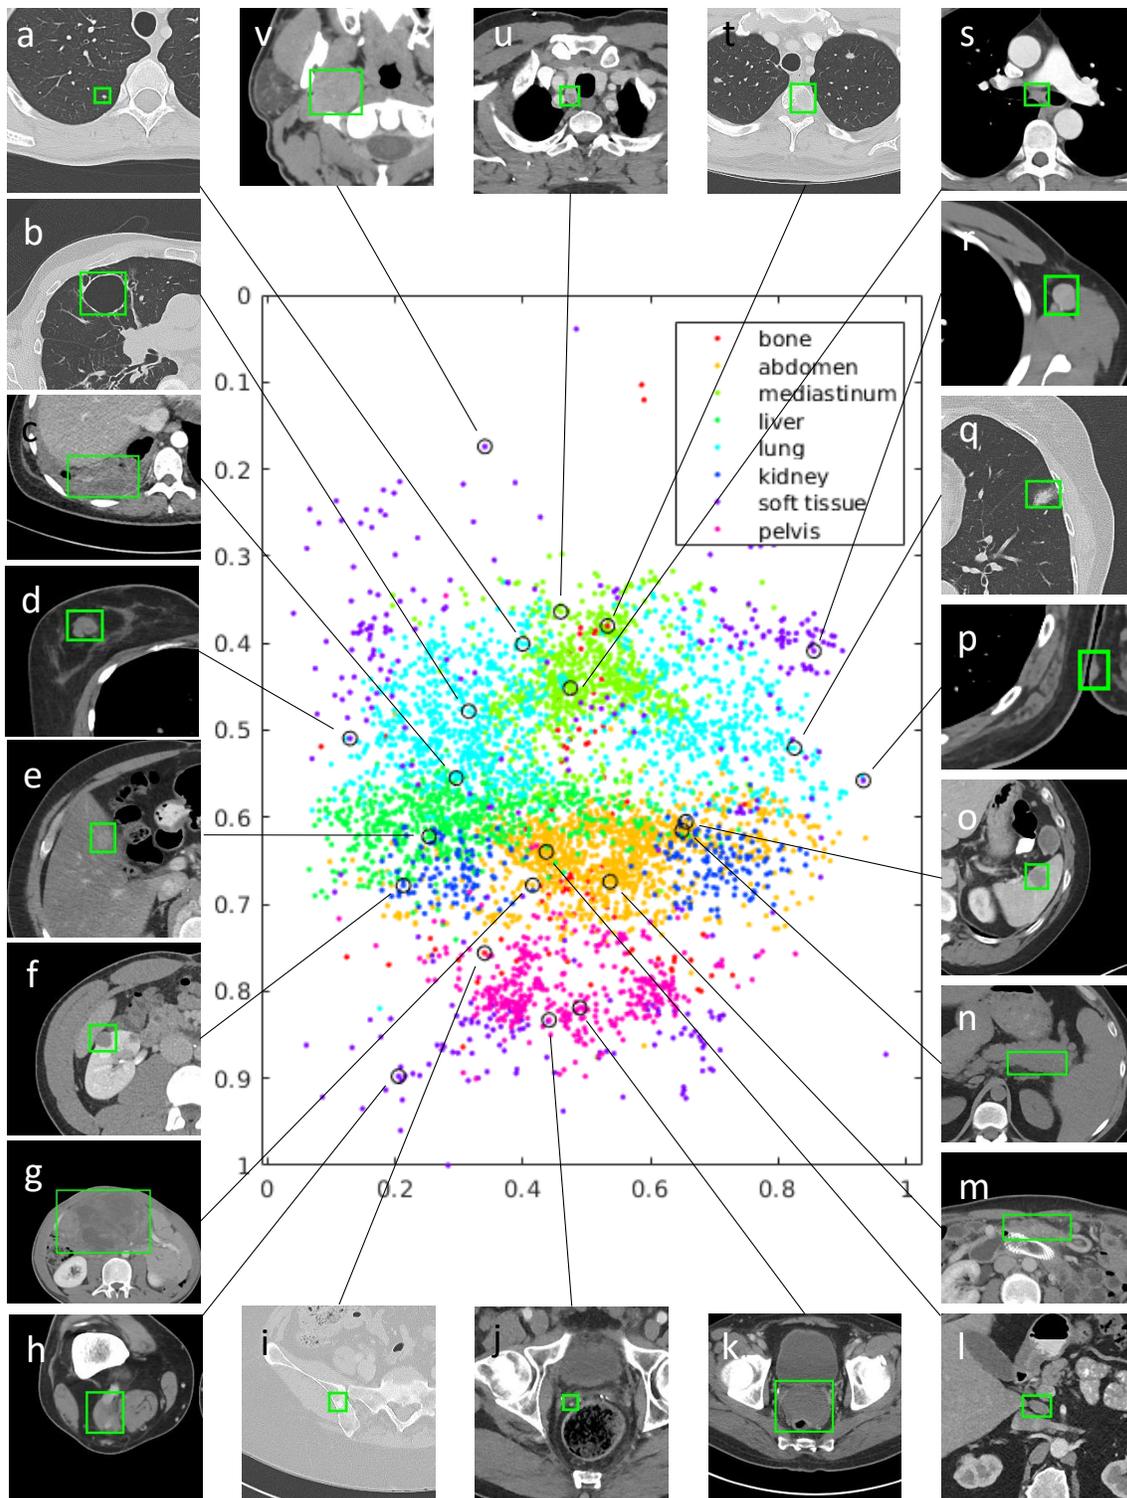


Figure 3. Visualization of the DeepLesion dataset (test set). The x - and y -axes of the scatter map correspond to the x - and z -coordinates of the relative body location of each lesion, respectively. Therefore, this map is similar to a frontal view of the human body. Colors indicate the manually labeled lesion types. Sample lesions are exhibited to show the great diversity of DeepLesion, including: a. lung nodule; b. lung cyst; c. costophrenic sulcus (lung) mass/fluid; d. breast mass; e. liver lesion; f. renal mass; g. large abdominal mass; h. posterior thigh mass; i. iliac sclerotic lesion; j. perirectal lymph node (LN); k. pelvic mass; l. periportal LN; m. omental mass; n. peripancreatic lesion; o. splenic lesion; p. subcutaneous/skin nodule; q. ground glass opacity; r. axillary LN; s. subcarinal LN; t. vertebral body metastasis; u. thyroid nodule; v. neck mass.

2. Self-Supervised Body-Part Regressor: Implementation Details

To train SSBR, we randomly pick 800 unlabeled CT volumes of 420 subjects from DeepLesion. Each axial slice in the volumes is resized to 128×128 pixels. No further preprocessing or data augmentation was performed. In each mini-batch, we randomly select 256 slices from 32 volumes (8 equidistant slices in each volume, see Eq. 2 in the paper) for training. The network is trained using stochastic gradient descent with a learning rate of 0.002. It generally converges in 1.5K iterations.

The sample lesions in Fig. 3 can be used to qualitatively evaluate the learned slice scores, or z -coordinates. We also conducted a preliminary experiment to quantitatively assess SSBR. A test set including 18,195 slices subsampled from 260 volumes of 140 new subjects are collected. They are manually labeled as one of the 3 classes: chest (5903 slices), abdomen (6744), or pelvis (5548). The abdomen class starts from the upper border of the liver and ends at the upper border of the ilium. Then, we set two thresholds on the slice scores to classify them to the three classes. The classification accuracy is 95.99%, with all classification errors appearing at transition regions (chest-abdomen, abdomen-pelvis) partially because of their ambiguity. The result proves the effectiveness of SSBR. More importantly, SSBR is trained on unlabeled volumes that are abundant in every hospital's database, thus zero annotation effort is needed.

3. Content-based Lesion Retrieval: More Results

More examples of lesion retrieval are shown in Fig. 4. We try to exhibit typical examples of all lesion types. The last row is a failure case. Most retrieved lesions are similar with the query ones in type, location, and size. More importantly, most retrieved lesions and the query ones come from semantically similar body structures that are not specified in the training labels. The failure cases in Fig. 4 have dissimilar types with the query ones. They were retrieved mainly because they have similar location, size, and appearance with query ones.

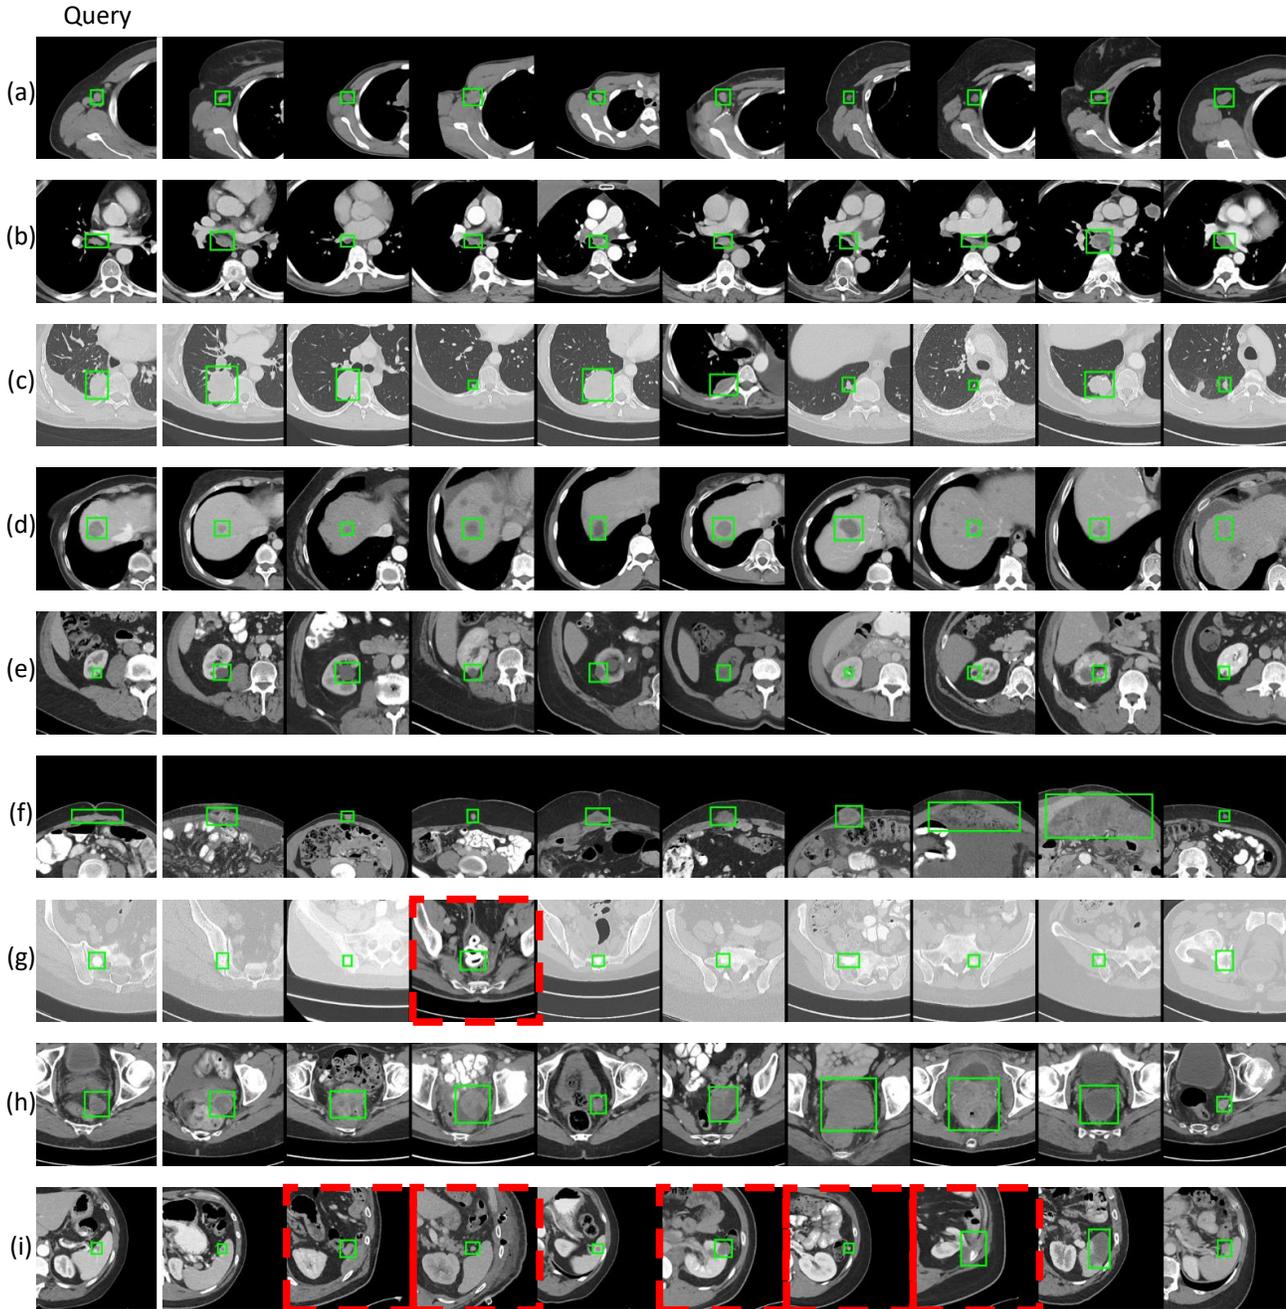


Figure 4. More examples of query lesions (first column) and the top-9 retrieved lesions on the test set of DeepLesion. We constrain that the query and all retrieved lesions must come from different patients. Red dashed boxes indicate incorrect results. The lesions in each row are: (a) Right axillary lymph nodes; (b) subcarinal lymph nodes; (c) lung masses or nodules near the pleura; (d) liver lesions near the liver dome; (e) right kidney lesions; (f) lesions near the anterior abdomen wall; (g) lesions on pelvic bones except the one in the red box, which is a peripherally calcified mass. (h) inferior pelvic lesions; (i) spleen lesions except the ones in red boxes.

4. Intra-Patient Lesion Matching: An Example

To provide an intuitive illustration of the lesion matching task, we show lesions of a sample patient in Fig. 5, with their

lesion graph in Fig. 6 and the final extracted lesion sequences in Fig. 7. We show that the lesion graph and Algo. 1 in the paper can be used to accurately match lesions in multiple studies.

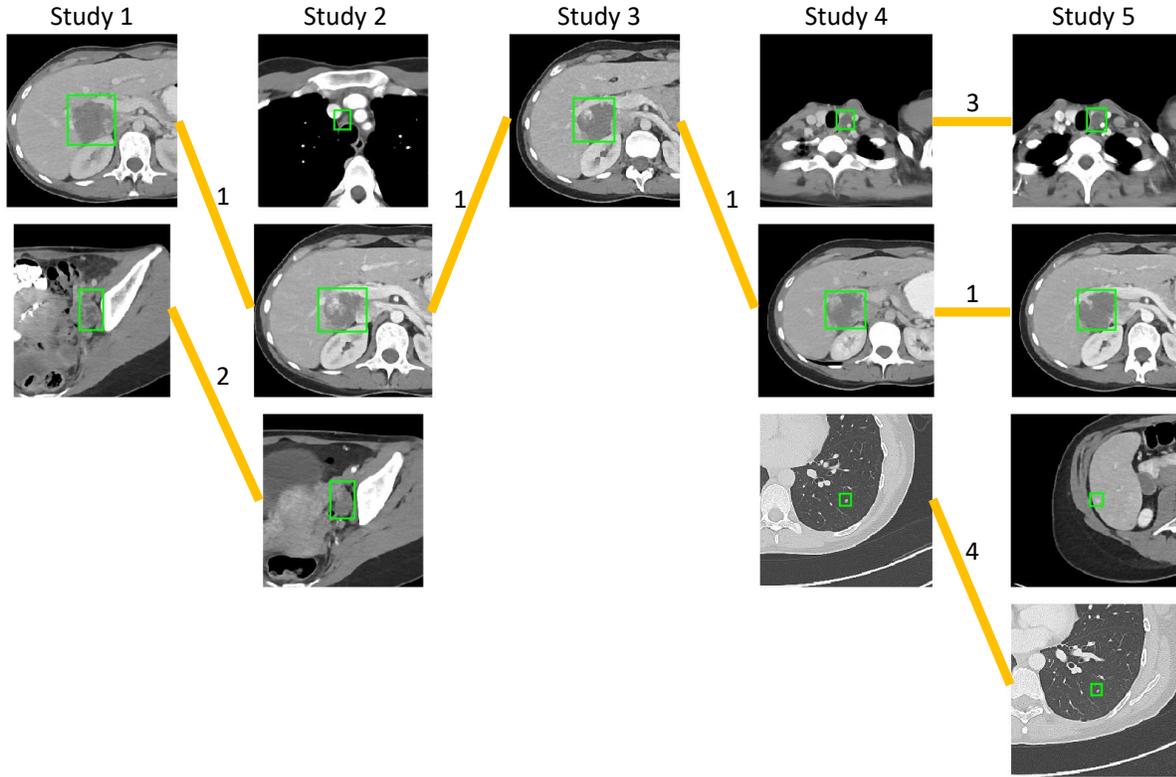


Figure 5. All lesions of a sample patient in DeepLesion. Lesions in each study (CT examination) are listed in a column. Not all lesions occur in each study, because the scan ranges of each study vary and radiologists only mark a few target lesions. We group the same lesion instances to sequences. Four sequences are found and marked in the figure, where the numbers on the connections represent the lesion IDs.

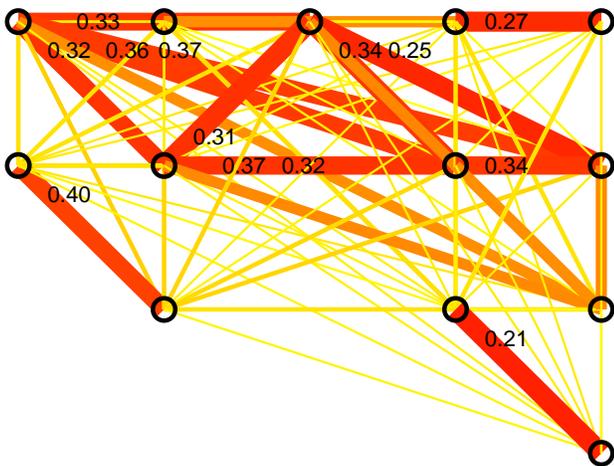


Figure 6. The intra-patient lesion graph of the patient in Fig. 5. For clarity, the lesions in Fig. 5 are replaced by nodes in this figure. The numbers on the edges are the Euclidean distances between nodes. We only show small distances in the figure. Red, thick edges indicate smaller distances. Note that some edges may overlap with other edges or nodes.

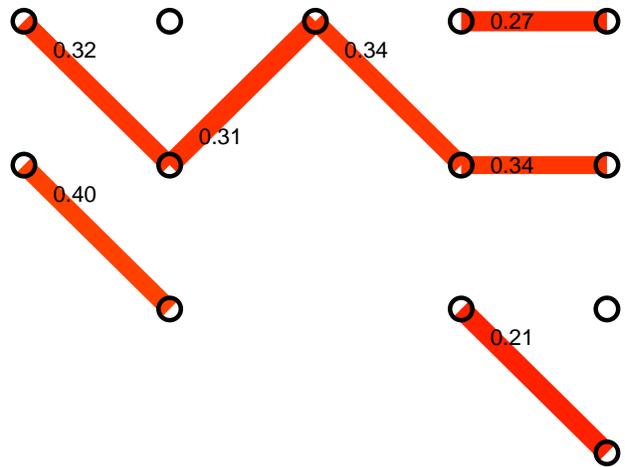


Figure 7. The final lesion sequences found by processing the lesion graph in Fig. 6 using Algo. 1 in the paper. They are the same with the ground-truth in Fig. 5.