



Training Multi-organ Segmentation Networks with Sample Selection by Relaxed Upper Confident Bound

Yan Wang¹(✉), Yuyin Zhou¹, Peng Tang², Wei Shen^{1,3}, Elliot K. Fishman⁴,
and Alan L. Yuille¹

¹ Johns Hopkins University, Baltimore, USA
ywang372@jhu.edu

² Huazhong University of Science and Technology, Wuhan, China

³ Shanghai University, Shanghai, China

⁴ Johns Hopkins University School of Medicine, Baltimore, USA

Abstract. Convolutional neural networks (CNNs), especially fully convolutional networks, have been widely applied to automatic medical image segmentation problems, e.g., multi-organ segmentation. Existing CNN-based segmentation methods mainly focus on looking for increasingly powerful network architectures, but pay less attention to data sampling strategies for training networks more effectively. In this paper, we present a simple but effective sample selection method for training multi-organ segmentation networks. Sample selection exhibits an exploitation-exploration strategy, i.e., exploiting hard samples and exploring less frequently visited samples. Based on the fact that very hard samples might have annotation errors, we propose a new sample selection policy, named Relaxed Upper Confident Bound (RUCB). Compared with other sample selection policies, e.g., Upper Confident Bound (UCB), it exploits a range of hard samples rather than being stuck with a small set of very hard ones, which mitigates the influence of annotation errors during training. We apply this new sample selection policy to training a multi-organ segmentation network on a dataset containing 120 abdominal CT scans and show that it boosts segmentation performance significantly.

1 Introduction

The field of medical image segmentation has made significant advances riding on the wave of deep convolutional neural networks (CNNs). Training convolutional deep networks (CNNs), especially fully convolutional networks (FCNs) [6], to automatically segment organs from medical images, such as CT scans, has become the dominant method, due to its outstanding segmentation performance. It also sheds lights to many clinical applications, such as diabetes inspection, organic cancer diagnosis, and surgical planning.

To approach human expert performance, existing CNN-based segmentation methods mainly focus on looking for increasingly powerful network architectures,




Fig. 1. Examples in a abdominal CT scans dataset which have annotations errors. Left: vein is included in pancreas segmentation; Middle & Right: missing pancreas header.

e.g., from plain networks to residual networks [5, 10], from single stage networks to cascaded networks [13, 16], from networks with a single output to networks with multiple side outputs [8, 13]. However, there is much less study of how to select training samples from a fixed dataset to boost performance.

In the training procedure of current state-of-the-art CNN-based segmentation methods [4, 11, 12, 17], training samples (2D slices for 2D FCNs and 3D sub-volumes for 3D FCNs) are randomly selected to iteratively update network parameters. However, some samples are much harder to segment than others, e.g., those which contain more organs with indistinct boundaries or with small sizes. It is known that using hard sample selection, or called bootstrapping¹, for training deep networks yields faster training, higher accuracy, or both [7, 14, 15]. Hard sample selection strategies for object detection [14] and classification [7, 15] base their selection on the training loss for each sample, but some samples are hard due to annotation errors, as shown in Fig. 1. This problem may not be significant for the tasks in natural images, but for the tasks in medical images, such as multi-organ segmentation, which usually require very high accuracy, and thus the influence of annotation errors is more significant. Our experiments show that the training losses of samples (such as the samples in Fig. 1) with annotation errors are very large, and even larger than real hard samples.

To address this problem, we propose a new hard sample selection policy, named Relaxed Upper Confident Bound (RUCB). Upper Confident Bound (UCB) [2] is a classic policy to deal with *exploitation-exploration* trade-offs [1], e.g., exploiting hard samples and exploring less frequently visited samples for sample selection. UCB was used for object detection in natural images [3], but UCB is easy to be stuck with some samples with very large losses, as the selection procedure goes on. In our RUCB, we relax this policy by selecting hard samples from a larger range, but with higher probability for harder samples, rather than only selecting some very hard samples as the selection procedure goes on. RUCB can escape from being stuck with a small set of very hard samples, which can mitigate the influence of annotation errors. Experimental results on a dataset containing 120 abdominal CT scans show that the proposed Relaxed Upper Confident Bound policy boosts multi-organ segmentation performance significantly.

¹ In this paper, we only consider the bootstrapping procedure that selects samples from a fixed dataset.

2 Methodology

Given a 3D CT scan $V = (v_j, j = 1, \dots, |V|)$, the goal of multi-organ segmentation is to predict the label of all voxels in the CT scan $\hat{Y} = (\hat{y}_j, j = 1, \dots, |V|)$, where $\hat{y}_j \in \{0, 1, \dots, |\mathcal{L}|\}$ denotes the predicted label for each voxel v_j , i.e., if v_j is predicted as a background voxel, then $\hat{y}_j = 0$; and if v_j is predicted as an organ in the organ space \mathcal{L} , then $\hat{y}_j = 1, \dots, |\mathcal{L}|$. In this section, we first review the basics of the Upper Confident Bound policy [2], then elaborate our proposed Relaxed Upper Confident Bound policy on sample selection for multi-organ segmentation.

2.1 Upper Confident Bound (UCB)

The Upper Confident Bound (UCB) [2] policy is widely used to deal with the exploration versus exploitation dilemma, which arises in the multi-armed bandit (MAB) problem [9]. In a K -armed bandit problem, each arm $k = 1, \dots, K$ is recorded by an unknown distribution associated with an unknown expectation. In each trial $t = 1, \dots, T$, a learner takes an action to choose one of K alternatives $g(t) \in \{1, \dots, K\}$ and collects a reward $x_{g(t)}^{(t)}$. The objective of this problem is to maximize the long-run cumulative expected reward $\sum_{t=1}^T x_{g(t)}^{(t)}$. But, as the expectations are unknown, the learner can only make a judgement based on the record of the past trails.

At trial t , the UCB selects the alternative k maximizing $\bar{x}_k + \sqrt{\frac{2 \ln n}{n_k}}$, where $\bar{x}_k = \sum_{t=1}^n x_k^{(t)} / n_k$ is the average reward obtained from the alternative k based on the previous trails, $x_k^{(t)} = 0$ if x_k is not chosen in the t -th trail. n_k is the number of times alternative k has been selected so far and n is the total number of trail done. The first term is the exploitation term, whose value is higher if the expected reward is larger; and the second term is the exploration term, which grows with the total number of actions that have been taken but shrinks with the number of times this particular action has been tried. At the beginning of the process, the exploration term dominates the selection, but as the selection procedure goes on, the one with the best expected reward will be chosen.

2.2 Relaxed Upper Confident Bound (RUCB) Bostrapping

Fully convolutional networks (FCNs) [6] are the most popular model for multi-organ segmentation. In a typical training procedure of an FCN, a sample (e.g., a 2D slice) is randomly selected in each iteration to calculate the model error and update model parameters. To train this FCN more effectively, a better strategy is to use hard sample selection, rather than random sample selection. As sample selection exhibits an exploitation-exploration trade-off, i.e., exploiting hard samples and exploring less frequently visited samples, we can directly apply UCB to select samples, where the reward of a sample is defined as the network loss function w.r.t. it. However, as the selection procedure goes on, only a small set of samples with the very large reward will be selected for next iteration according

to UCB. The selected sample may not be a proper hard sample, but a sample with annotation errors, which inevitably exist in medical image data as well as other image data. Next, we introduce our Relaxed Upper Confident Bound (RUCB) policy to address this issue.

Procedure. We consider that training an FCN for multi-organ segmentation, where the input images are 2D slices from axial directions. Given a training set $\mathcal{S} = \{(\mathbf{I}_i, \mathbf{Y}_i)\}_{i=1}^M$, where \mathbf{I}_i and \mathbf{Y}_i denote a 2D slice and its corresponding label map, and M is the number of the 2D slices, like the MAB problem, each slice \mathbf{I}_i is set to be associated with the number of times it was selected n_i and the average reward obtained through the training \bar{J}_i . After training an initial FCN with randomly sampling slices from the training set, it is bootstrapped several times by sampling hard and less frequently visited slices. In the sample selection procedure, rewards are assigned to each training slice once, then the next slice to train FCN is chosen by the proposed RUCB. The reward of this slice is fed into RUCB and the statistics in RUCB are updated. This process is then repeated to select another slice based on the updated statistics, until a max-iteration N is reached. Statistics are reset to 0 before beginning a new bootstrapping phase since slices that are chosen in previous rounds may no longer be informative.

Relaxed Upper Confident Bound. We denote the corresponding label map of the input 2D slice $\mathbf{I}_i \subset \mathbb{R}^{H \times W}$ as $\mathbf{Y}_i = \{y_{i,j}\}_{j=1, \dots, H \times W}$. If \mathbf{I}_i is selected to update the FCN in the t -th iteration, the reward obtained for \mathbf{I}_i is computed by

$$\mathcal{J}_i^{(t)}(\Theta) = -\frac{1}{H \times W} \left[\sum_{j=1}^{H \times W} \sum_{l=0}^{|\mathcal{L}|} \mathbf{1}(y_{i,j} = l) \log p_{i,j,l}^{(t)} \right], \quad (1)$$

where $p_{i,j,l}^{(t)}$ is the probability that the label of the j -th pixel in the input slice is l , and $p_{i,j,l}^{(t)}$ is parameterized by the network parameter Θ . If \mathbf{I}_i is not selected to update the FCN in the t -th iteration, $\mathcal{J}_i^{(t)}(\Theta) = 0$. After n iterations, the next slice to be selected by UCB is the one maximizing $\bar{J}_i^{(n)} + \sqrt{2 \ln n / n_i}$, where $\bar{J}_i^{(n)} = \sum_{t=1}^n \mathcal{J}_i^{(t)}(\Theta) / n_i$.

Preliminary experiments show that reward defined above is usually around $[0, 0.35]$. The exploration term dominates the exploitation term. We thus normalize the reward to make a balance between exploitation and exploration by

$$\tilde{J}_i^{(n)} = \min \left\{ \beta, \frac{\beta}{2} \frac{\bar{J}_i^{(n)}}{\sum_{i=1}^M \bar{J}_i^{(n)} / M} \right\}, \quad (2)$$

where the min operation ensures that the score lies in $[0, \beta]$. Then the UCB score for \mathbf{I}_i is calculated as

$$q_i^{(n)} = \tilde{J}_i^{(n)} + \sqrt{\frac{2 \ln n}{n_i}}. \quad (3)$$

Algorithm 1. Relaxed Upper Confident Bound

Input : FCN parameter Θ , input training slices $\{\mathbf{I}_i\}_{i=1,\dots,M}$;
parameters α and β , max number of iterations T ;

Output: FCN parameter Θ ;

- 1 total number of times slices are selected $n \leftarrow 0$;
- 2 number of times slice $\mathbf{I}_1, \dots, \mathbf{I}_m$ are selected $n_1, \dots, n_m \leftarrow 0$;
- 3 running index $i \leftarrow 0, \mathcal{J}_1^{(1)}, \dots, \mathcal{J}_M^{(M)} \leftarrow 0$;
- 4 **repeat**
- 5 $i \leftarrow i + 1, n_i \leftarrow n_i + 1, n \leftarrow n + 1$;
- 6 Compute $\mathcal{J}_i^{(i)}$ by Eq. 1;
- 7 $\bar{\mathcal{J}}_i^{(M)} = \sum_{t=1}^M \mathcal{J}_i^{(i)} / n_i$;
- 8 **until** $n = M$;
- 9 $\forall i$, compute $\tilde{\mathcal{J}}_i^{(M)}$ by Eq. 2, compute $q_i^{(M)}$ by Eq. 3;
- 10 $\mu = \sum_{i=1}^M q_i^{(M)} / M, \sigma = \text{std}(q_i^{(M)})$; iteration $t \leftarrow 0$;
- 11 **repeat**
- 12 $t \leftarrow t + 1, \alpha \sim \mathcal{U}(0, a)$;
- 13 $K = \sum_{i=1}^M (\mathbf{1}(q_i^{(M)} > \mu + \alpha\sigma))$;
- 14 randomly select a slice \mathbf{I}_i from the set $\{\mathbf{I}_i | q_i^{(n)} \in \mathcal{D}_K(\{q_i^{(n)}\}_{i=1}^M)\}$;
- 15 $n_i \leftarrow n_i + 1, n \leftarrow n + 1$;
- 16 Compute $\mathcal{J}_i^{(t)}$ by Eq. 1, $\Theta \leftarrow \arg \min_{\Theta} \mathcal{J}_i^{(t)}(\Theta)$;
- 17 $\bar{\mathcal{J}}_i^{(n)} = \sum_{t=1}^n \mathcal{J}_i^{(t)} / n_i$;
- 18 $\forall i$, compute $\tilde{\mathcal{J}}_i$ by Eq. 2, compute $q_i^{(n)}$ by Eq. 3;
- 19 **until** $t = T$;

As the selection procedure goes on, the exploitation term of Eq. 3 will dominate the selection, i.e., only some very hard samples will be selected. But, these hard samples may have annotation errors. In order to alleviate the influence of annotation errors, we propose to introduce more randomness in UCB scores to relax the largest loss policy. After training an initial FCN with randomly sampling slices from the training set, we assign an initial UCB score $q_i^{(M)} = \tilde{\mathcal{J}}_i^{(M)} + \sqrt{2 \ln M / I}$ to each slice \mathbf{I}_i in the training set. Assume the UCB scores of all samples follow a normal distribution $\mathcal{N}(\mu, \sigma)$. Hard samples are regarded as slices whose initial UCB scores are larger than μ . Note that initial UCB scores are only decided by the exploitation term. In each iteration of our bootstrapping procedure, we count the number of samples that lie in the range $[\mu + \alpha \cdot \text{std}(q_i^{(M)}), +\infty]$, denoted by K , where α is drawn from a uniform distribution $[0, a]$ ($a = 3$ in our experiment), then a sample is selected randomly from the set $\{\mathbf{I}_i | q_i^{(n)} \in \mathcal{D}_K(\{q_i^{(n)}\}_{i=1}^M)\}$ to update the FCN, where $\mathcal{D}_K(\cdot)$ denote the K largest values in a set. Here we count the number of hard samples according to a dynamic range, because we do not know the exact range of hard samples. This dynamic region enables our bootstrapping to select hard samples from a larger range with higher probability for harder samples, rather than only selecting some very hard samples. We name our sample selection policy Relaxed Upper Confident Bound (RUCB), as we choose hard samples in a larger range, which introduces more variance to the hard samples. The training procedure for RUCB is summarized in Algorithm 1.

3 Experimental Results

3.1 Experimental Setup

Dataset: We evaluated our algorithm on 120 abdominal CT scans of normal cases under IRB (Institutional Review Board) approved protocol. CT scans are contrast enhanced images in portal venous phase, obtained by Siemens SOMATOM Sensation64 and Definition CT scanners, composed of (319–1051) slices of (512×512) images, and have voxel spatial resolution of $([0.523 - 0.977] \times [0.523 - 0.977] \times 0.5) \text{ mm}^3$. Sixteen organs (including aorta, celiac AA, colon, duodenum, gallbladder, interior vena cava, left kidney, right kidney, liver, pancreas, superior mesenteric artery, small bowel, spleen, stomach, and large veins) were segmented by four full-time radiologists, and confirmed by an expert. This dataset is a high quality dataset, but a small portion of error is inevitable, as shown in Fig. 1. Following the standard cross-validation strategy, we randomly partition the dataset into four complementary folds, each of which contains 30 CT scans. All experiments are conducted by four-fold cross-validation, i.e., training the models on three folds and testing them on the remaining one, until four rounds of cross-validation are performed using different partitions.

Evaluation Metric: The performance of multi-organ segmentation is evaluated in terms of Dice-Sørensen similarity coefficient (DSC) over the whole CT scan. We report the average DSC score together with the standard deviation over all testing cases.

Implementation Details: We use FCN-8s model [6] pre-trained on PascalVOC in caffe toolbox. The learning rate is fixed to be 1×10^{-9} and all the networks are trained for 80K iterations by SGD. The same parameter setting is used for all sampling strategies. Three bootstrapping phases are conducted, at 20,000, 40,000 and 60,000 respectively, i.e., the max number of iterations for each bootstrapping phase is $T = 20,000$. We set $\beta = 2$, since $\sqrt{2 \ln n/n_i}$ is in the range of [3.0, 5.0] in bootstrapping phases.

3.2 Evaluation of RUCB

We evaluate the performance of the proposed sampling algorithm (RUCB) with other competitors. Three sampling strategies considered for comparisons are (1) uniform sampling (Uniform); (2) online hard example mining (OHEM) [14]; and (3) using UCB policy (i.e., select the slice with the largest UCB score during each iteration) in bootstrapping.

Table 1 summarizes the results for 16 organs. Experiments show that images with wrong annotations are with large rewards, even larger than real hard samples after training an initial FCN. The proposed RUCB outperforms over all baseline algorithms in terms of average DSC. We see that RUCB achieves much better performance for organs such as *Adrenal gland* (from 29.33% to 36.76%),

Table 1. DSC (%) of sixteen segmented organs (mean \pm standard deviation).

Organs	Uniform	OHEM	UCB	RUCB (ours)
Aorta	81.53 \pm 4.50	77.49 \pm 5.90	81.02 \pm 4.50	81.03 \pm 4.40
Adrenal gland	29.33 \pm 16.26	31.44 \pm 16.71	33.75 \pm 16.26	36.76 \pm 17.28
Celiac AA	34.49 \pm 12.92	33.34 \pm 13.86	35.89 \pm 12.92	38.45 \pm 12.53
Colon	77.51 \pm 7.89	73.20 \pm 8.94	76.40 \pm 7.89	77.56 \pm 8.65
Duodenum	63.39 \pm 12.62	59.68 \pm 12.32	63.10 \pm 12.62	64.86 \pm 12.18
Gallbladder	79.43 \pm 23.77	77.82 \pm 23.58	79.10 \pm 23.77	79.68 \pm 23.46
IVC	78.75 \pm 6.54	73.73 \pm 8.59	77.10 \pm 6.54	78.57 \pm 6.69
Left kidney	95.35 \pm 2.53	94.24 \pm 8.95	95.53 \pm 2.53	95.57 \pm 2.29
Right kidney	94.48 \pm 9.49	94.23 \pm 9.19	94.39 \pm 9.49	95.40 \pm 3.62
Liver	96.03 \pm 1.70	90.43 \pm 4.74	95.68 \pm 1.70	96.00 \pm 1.28
Pancreas	77.86 \pm 9.92	75.32 \pm 10.42	78.25 \pm 9.92	78.48 \pm 9.86
SMA	45.36 \pm 14.36	47.18 \pm 12.75	44.63 \pm 14.36	49.59 \pm 13.62
Small bowel	72.35 \pm 13.30	67.44 \pm 13.22	72.16 \pm 13.30	72.88 \pm 13.98
Spleen	95.32 \pm 2.17	94.56 \pm 2.41	95.16 \pm 2.17	95.09 \pm 2.44
Stomach	90.62 \pm 6.51	86.37 \pm 8.53	90.70 \pm 6.51	90.92 \pm 5.62
Veins	64.95 \pm 19.96	60.87 \pm 19.02	62.70 \pm 19.96	65.13 \pm 20.15
AVG	73.55 \pm 10.28	71.08 \pm 11.20	73.47 \pm 10.52	74.75 \pm 9.88

Celiac AA (34.49% to 38.45%), *Duodenum* (63.39% to 64.86%), *Right kidney* (94.48% to 95.40%), *Pancreas* (77.86% to 78.48%) and *SMA* (45.36% to 49.59%), compared with Uniform. Most of the organs listed above are small organs which are difficult to segment, even for radiologists, and thus they may have more annotation errors.

OHEM performs worse than Uniform, suggesting that directly sampling among slices with largest average rewards during bootstrapping phase cannot help to train a better FCN. UCB obtains even slightly worse DSC compared with Uniform, as it only focuses on some hard examples which may have errors.

To better understand UCB and RUCB, some of the hard samples selected more frequently are shown in Fig. 2. Some slices selected by UCB contain obvious errors such as *Colon* annotation for the first one. Slices selected by RUCB are very hard to segment since it contains many organs including very small ones.

Parameter Analysis. α is an important hyper-parameter for our RUCB. We vary it in the following range: $\alpha \in \{0, 1, 2, 3\}$, to see how the performance of some organs changes. The DSCs of *Adrenal gland* and *Celiac AA* are 35.36 ± 17.49 and 38.07 ± 12.75 , 32.27 ± 16.25 and 36.97 ± 12.92 , 34.42 ± 17.17 and 36.68 ± 13.73 , 32.65 ± 17.26 and 37.09 ± 12.15 , respectively. Using a fixed α , the performance decreases. We also test the results when K is a constant number, i.e., $K = 5000$. The DSC of *Adrenal gland* and *Celiac AA* are 33.55 ± 17.02 and 36.80 ± 12.91 . Compared with UCB, the results further verify that relaxing the UCB score can boost the performance.

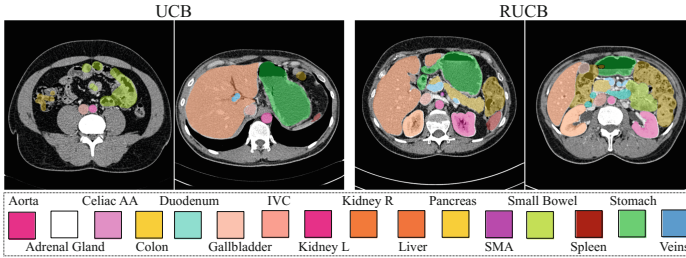


Fig. 2. Visualization of samples selected frequently by left: UCB and right: RUCB. Ground-truth annotations are marked in different colors.

4 Conclusion

We proposed Relaxed Upper Confident Bound policy for sample selection in training multi-organ segmentation networks, in which the exploitation-exploration trade-off is reflected on one hand by the necessity for trying all samples to train a basic classifier, and on the other hand by the demand of assembling hard samples to improve the classifier. It exploits a range of hard samples rather than being stuck with a small set of very hard samples, which mitigates the influence of annotation errors during training. Experimental results showed the effectiveness of the proposed RUCB sample selection policy. Our method can be also used for training 3D patch-based networks, and with other modality medical images.

Acknowledgement. This work was supported by the Lustgarten Foundation for Pancreatic Cancer Research and also supported by NSFC No. 61672336. We thank Prof. Seyoun Park and Dr. Lingxi Xie for instructive discussions.

References

1. Auer, P.: Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.* **3**, 397–422 (2002)
2. Auer, P., Cesa-Bianchi, N., Fischer, P.: Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.* **47**(2), 235–256 (2002)
3. Canévet, O., Fleuret, F.: Large scale hard sample mining with monte carlo tree search. In: *Proceedings of the CVPR*, pp. 5128–5137 (2016)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3D U-Net: learning dense volumetric segmentation from sparse annotation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) *MICCAI 2016*. LNCS, vol. 9901, pp. 424–432. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_49
5. Fakhry, A., Zeng, T., Ji, S.: Residual deconvolutional networks for brain electron microscopy image segmentation. *IEEE Trans. Med. Imaging* **36**(2), 447–456 (2017)
6. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Computer Vision and Pattern Recognition* (2015)

7. Loshchilov, I., Hutter, F.: Online batch selection for faster training of neural networks. CoRR abs/1511.06343 (2015)
8. Merkow, J., Marsden, A., Kriegman, D., Tu, Z.: Dense volume-to-volume vascular boundary detection. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9902, pp. 371–379. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46726-9_43
9. Robbins, H.: Some aspects of the sequential design of experiments. Bull. Am. Math. Soc. **58**(5), 527–535 (1952)
10. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
11. Roth, H., et al.: Hierarchical 3D fully convolutional networks for multi-organ segmentation. CoRR abs/1704.06382 (2017)
12. Roth, H.R., Lu, L., Farag, A., Sohn, A., Summers, R.M.: Spatial aggregation of holistically-nested networks for automated pancreas segmentation. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9901, pp. 451–459. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46723-8_52
13. Shen, W., Wang, B., Jiang, Y., Wang, Y., Yuille, A.L.: Multi-stage multi-recursive-input fully convolutional networks for neuronal boundary detection. In: Proceedings of the ICCV, pp. 2410–2419 (2017)
14. Shrivastava, A., Gupta, A., Girshick, R.B.: Training region-based object detectors with online hard example mining. In: CVPR, pp. 761–769 (2016)
15. Simo-Serra, E., Trulls, E., Ferraz, L., Kokkinos, I., Moreno-Noguer, F.: Fracking deep convolutional image descriptors. CoRR abs/1412.6537 (2014)
16. Wang, Y., Zhou, Y., Shen, W., Park, S., Fishman, E.K., Yuille, A.L.: Abdominal multi-organ segmentation with organ-attention networks and statistical fusion. CoRR abs/1804.08414 (2018)
17. Zhou, Y., Xie, L., Shen, W., Wang, Y., Fishman, E.K., Yuille, A.L.: A fixed-point model for pancreas segmentation in abdominal CT scans. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10433, pp. 693–701. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66182-7_79