

Boundary-RL: Reinforcement Learning for Weakly-Supervised Prostate Segmentation in TRUS Images

Weixi Yi¹, Vasilis Stavrinides², Zachary M.C. Baum¹, Qianye Yang¹, Dean C. Barratt¹, Matthew J. Clarkson¹, Yipeng Hu¹, Shaheer U. Saeed¹

¹ Wellcome / EPSRC Centre for Interventional and Surgical Sciences; Centre for Medical Image Computing; and Department of Medical Physics and Biomedical Engineering, University College London, London, UK

² Division of Surgery & Interventional Science, University College London; and Department of Urology, UCL Hospital NHS Foundation Trust, London, UK
`weixi.yi.22@ucl.ac.uk`

Abstract. We propose Boundary-RL, a novel weakly supervised segmentation method that utilises only patch-level labels for training. We envision segmentation as a boundary detection problem, rather than a pixel-level classification as in previous works. This outlook on segmentation may allow for boundary delineation under challenging scenarios such as where noise artefacts may be present within the region-of-interest (ROI) boundaries, where traditional pixel-level classification-based weakly supervised methods may not be able to effectively segment the ROI. Particularly of interest, ultrasound images, where intensity values represent acoustic impedance differences between boundaries, may also benefit from the boundary delineation approach. Our method uses reinforcement learning to train a controller function to localise boundaries of ROIs using a reward derived from a pre-trained boundary-presence classifier. The classifier indicates when an object boundary is encountered within a patch, serving as weak supervision, as the controller modifies the patch location in a sequential Markov decision process. The classifier itself is trained using only binary patch-level labels of object presence, the only labels used during training of the entire boundary delineation framework. The use of a controller ensures that sliding window over the entire image is not necessary and reduces possible false-positives or -negatives by minimising number of patches passed to the boundary-presence classifier. We evaluate our approach for a clinically relevant task of prostate gland segmentation on trans-rectal ultrasound images. We show improved performance compared to other tested weakly supervised methods, using the same labels e.g., multiple instance learning.

Keywords: TRUS · Weak Supervision · Reinforcement Learning.

1 Introduction

Automated segmentation of anatomical structures and other regions-of-interest (ROI) plays a vital role in the field of medical imaging [26, 18]. Accurate de-

lineation and contouring of ROI boundaries is an essential step in procedures such as treatment planning e.g., for radiotherapy [24], longitudinal diagnostic or prognostic planning e.g., by tracking fetal dimensions [35] or organ sizes such as prostate volume [32], for locating abnormalities within imaged structures e.g., for lung abnormality detection [5], and for trans-rectal ultrasound (TRUS) guided prostate biopsy or brachytherapy [12]. In particular, of interest in this work, TRUS-guided biopsy has recently been aided by automated segmentation of the gland boundaries for precise sampling along the gland for accurate diagnoses [12, 29, 14]. Recent advancements in the field have been driven by the introduction of fully supervised learning-based segmentation methods [30, 9] that require expert annotated data for learning. However, due to the time-consuming nature of pixel-level annotation of medical images, the requirement for specialized knowledge and significant inter-observer variability, the acquisition of ample annotated data for model training may be challenging [2, 4, 23, 1, 21].

Weakly-supervised semantic segmentation (WSSS) may address some of the above-mentioned issues with fully-supervised learning by employing weak labels, e.g., image-level classification labels, for learning segmentation [36, 10, 15]. However, segmenting the boundaries of the prostate from TRUS images with weak labels remains significantly challenging [7] due to: (1) often indistinct boundaries between prostate and surrounding tissues caused by the low contrast of soft tissues, (2) frequently missing segments of prostate boundaries due to the presence of shadowing artefacts, (3) considerable patient-dependent variability in prostate shape and size, (4) uneven intensity distribution of the prostate, and (5) lack of informative signals, e.g., in full supervision, where in commonly used approaches for weak supervision, presence of object to learn a valid prior such as global correlation between missing and distinct boundaries becomes challenging.

Within the literature on WSSS, methods based on multiple instance learning (MIL) have been proposed, with varying definitions of bags and patches used for training. For instance, MILinear [19], with its bag-splitting-based mechanism, is capable of iteratively generating new negative bags from positive ones. Xu et al. [33] integrated the concept of clustering and introduced a multi-cluster instance learning technique for segmenting cancerous and non-cancerous tissues. CAMEL [31] treated lattice patches as instances and histopathological images as bags. Increasing research efforts are oriented towards considering the image as a bag and treating each pixel within the image as an instance, for improved efficiency at inference, thereby transforming the weakly supervised segmentation task with only image-level annotations into an instance prediction task based on bag-level annotations. Jia et al. [8] proposed DWS-MIL, based on a fully convolutional network, for segmenting carcinogenic areas within histopathological images. Li et al. [13] introduced a self-attention mechanism to capture the correlation between instances within MIL. Additionally, some methods employ patch-level labels (i.e., binary labels of ROI presence within a patch or crop of a larger image) to achieve pixel-level segmentation with a comparatively lower annotation cost [6]. Most previous methods are application-specific with no general solutions.

We present a novel WSSS approach that only utilises patch-level binary classification labels of object-presence for training. These are more-time efficient to obtain than full segmentation maps traditionally used in fully-supervised learning. The WSSS in this work envisions the segmentation task as a boundary detection problem rather than a pixel-level classification, as in most previous works. The boundary detection follows a Markov Decision Process (MDP) whereby a controller function, trained using reinforcement learning (RL) based on reward function derived from a boundary presence classifier (similar to previously proposed task-based rewards [34, 22, 37, 3]), aims to localise the ROI boundary.

The use of RL for boundary delineation allows us to pose this as a sequential MDP whereby the patch-location refinement is done until the boundary of the object is found, as determined by the boundary presence classifier. As opposed to previously proposed WSSS methods that pose the segmentation task as a pixel-level classification, our method does not require a sliding window to be passed over the entire image to generate pixel-level labels for the segmentation map, which may be more efficient at inference, reducing the number of forward passes through the classifier. Rather, the boundary delineation may simply be generated using a trajectory that the controller follows within an image. Moreover, this approach of boundary delineation has the potential to perform well even when artefacts such as shadowing are present within the ROI but not occluding the boundary, where traditional weakly supervised pixel-wise classification approaches to segmentation may produce inaccurate results [16, 17, 27, 11, 19].

Contributions of this work are summarised: First, we propose a RL-based framework, Boundary-RL, trained only using patch-level labels of object presence, capable of effectively identifying object boundaries; Second, we evaluate Boundary-RL for a clinically relevant task of prostate gland boundary delineation on TRUS images using data from real prostate cancer patients; And lastly, we compare Boundary-RL to commonly used algorithms such as fully-supervised learning and MIL-based WSSS, which we adapt for our application, from previous work.

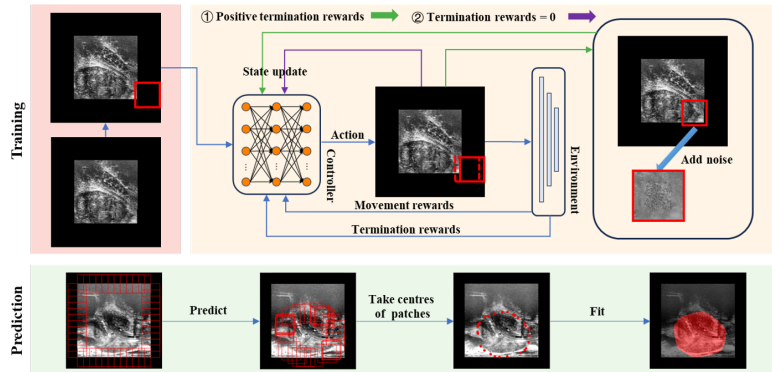


Fig. 1. The workflow of our proposed framework.

2 Methods

In our work, the boundary presence classifier, trained using only binary patch-level labels of ROI presence within the patch, serves as a weak signal that may be used to train the boundary detecting controller. The patch-localising controller ensures that the location of the patch is refined, starting from the edge of the image, until part of the ROI exists within the patch, as determined by the boundary presence classifier. This may be repeated with randomised starting locations in order to get an effective boundary delineation for the ROI.

2.1 Boundary presence classifier

The boundary presence classifier $f(\cdot; w) : \mathcal{X} \rightarrow \{0, 1\}$, with weights w , predicts object presence, given an image or image patch $x \in \mathcal{X}$, where \mathcal{X} is the domain of images or image patches. Here, 0, corresponds to object absence, and 1, corresponds to object presence. Image-label pairs $\{x_k, y_k\}_{k=1}^N$, where $y \in \{0, 1\}$, may be used to train the classifier f using a binary cross-entropy loss function by means of gradient descent following $w^* = \arg \min_w \mathbb{E}[L(y, f(x; w))]$ where $L(y, f(x; w))$ is the binary cross-entropy loss function and w^* are the optimal weights of the classifier which may be fixed after the initial pre-training stage. In practice the threshold for prediction probability is set as 0.9, so any classifier-predicted probability greater than 0.9 is rounded to 1, and 0 otherwise. The prediction probability of the classifier is treated as a boundary-presence probability prediction, which is used to inform the controller function, as outlined below.

2.2 Boundary-RL - delineating boundaries of objects using RL

We define an image patch x_c , of fixed dimensions, from image x as having coordinates $c = (i, j)$. The aim is for the controller function is to refine c until the ROI or its boundary are within the bounds of the patch.

The controller function $g(\cdot; \theta) : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{A}$, with weights θ , takes in the image $x \in \mathcal{X}$, the image patch $x_c \in \mathcal{X}$ and outputs an action $a \in \mathcal{A}$ where \mathcal{A} is the discrete action space defined as $a \in \mathcal{A} \in \{-1, 0, +1\}^2$. Here the two values for a correspond to the change in coordinates for the patch, in each axis i.e., $a = (\delta i, \delta j)$ which may be applied to the coordinates $c_t = (i, j)$ at time-step t , to translate the patch to the new coordinates $c_{t+1} = (i + \delta i, j + \delta j)$. In practice these actions are modelled as four discrete actions of up, down, left or right, that is to say that if $\delta i \neq 0$ then $\delta j = 0$ and if $\delta j \neq 0$ then $\delta i = 0$.

Interacting with an environment to learn boundary delineation The image x_t and pre-trained presence classifier $f(\cdot; w^*)$ may be considered part of a MDP environment with which the controller interacts, i.e., as a MDP. The environment may be defined as $(\mathcal{S}, \mathcal{A}, p, r, \pi, \gamma)$.

States: The observed state $s_t \in \mathcal{S}$ for this environment consists of the image x_t and the image patch localised by the controller $x_{c,t}$, i.e., $s_t = \{x_t, x_{c,t}\}$, where \mathcal{S} is the state space.

Actions: The action $a_t \in \mathcal{A}$, where \mathcal{A} is the action space, impacts the observed state at the subsequent time-step $t + 1$ as $a_t = (\delta i, \delta j)$, which leads to $c_{t+1} = (i + \delta i, j + \delta j)$, which in-turn results in an updated patch, $x_{c,t+1}$, which forms part of the updated state.

State transition distribution: The state-action conditioned state transition distribution is given by $p : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. Assuming next state s_{t+1} as input, given the current state s_t and action a_t , the probability of the next state is denoted as $p(s_{t+1} | s_t, a_t)$.

Rewards: The reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ takes the state-action pair as input and produces a reward $R_t = r(s_t, a_t)$, at time-step t . The reward consists of two components. 1) The first component is the movement reward r_{mov} , which is computed with the help of the estimated centroid of the prostate gland, given by coordinates $c_p = (i_p, j_p)$. The centroid of the prostate may be practically easy to obtain e.g., by means of pixel-intensity values or by manually labelling a subset of samples in roughly aligned images potentially utilising intensity-based registration. Then for time-step t , the L²-norm between the location of the patch and the centroid of the prostate gland is given by $\|c_t - c_p\|$. We then define $r_{mov} = +1$ if $\|c_t - c_p\| - \|c_{t-1} - c_p\| \leq 0$ and $r_{mov} = -1$ otherwise. 2) The second component is a termination reward r_{term} , given when the patch contains the ROI or its boundary, as determined by the boundary presence classifier function $f(\cdot; w^*)$, with pre-trained fixed weights w^* . We compute $r_{term} = f(x_{c,t}; w^*)$ for the termination reward, which is $f(x_{c,t}; w^*) = 1$ when the ROI is in the patch and $f(x_{c,t}; w^*) = 0$ otherwise. The two rewards may then be combined into our reward function $r(s_t, a_t) = r_{mov} + 100r_{term}$, where 100 is a scalar that allows weighting the two components. In this formulation, r_{mov} may be considered as reward shaping and r_{term} as the final sparse reward.

Policy: The probability of performing an action a_t is given by the policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, e.g., for time-step t the probability of performing action a_t is $\pi(a_t | s_t)$. Sampling an action according to a policy is denoted as $a_t \sim \pi(\cdot)$. Following the state transition distribution for sampling next states, the policy for sampling actions and the reward function for generating corresponding rewards, we can collect states, actions and rewards over multiple time-steps $(s_1, a_1, R_1, \dots, s_T, a_T, R_T)$, also called trajectories. These are used to update our controller.

Policy optimisation: The policy is modelled as a parametric neural network $\pi(\cdot; \theta)$ with weights θ . This neural network predicts parameters of a distribution from which to sample actions, which are the mean and standard deviation of a Gaussian distribution for continuous actions in this work [25]. The action is sampled from this predicted distribution $a \sim \pi(\cdot)$. We compute a cumulative reward $Q^{\pi(\cdot; \theta)}(s_t, a_t) = \sum_{k=0}^T \gamma^k R_{t+k}$ using a discount factor γ for future rewards. We can then use gradient ascent to obtain optimal policy parameters: $\theta^* = \arg \max_{\theta} \mathbb{E}[Q^{\pi(\cdot; \theta)}(s_t, a_t)]$.

Episodic training and terminal signals: In practice, we collect trajectories by defining episode termination as finding a patch containing the organ boundary. More concretely, we terminate an episode when $f(x_{c,t}; w^*) = 1$, we may denote this time-step at termination as t_{term} . At t_{term} , we add random Gaussian noise over our patch within the image such that finding the same patch-location on the next episode within the same image does not lead to $f(x_{c,t}; w^*) = 1$. This encourages finding unique patches as boundary points for the ROI. We conduct M episodes for each image until M patches have been found as the object boundary. At each episode the patch location is randomly selected as an edge-pixel on one of the four edges, so the controller learns to move from the edge of the image to the boundary of the object within the image. In practice, we also terminate the episode after T time-steps in case the boundary cannot be found after T steps, where in our formulation, $T = 1000$.

Using controller-predicted patches for object boundary delineation: The centre-points of all M patches are taken to be the estimated boundary points for the object, outliers are removed via a simple distance-based outlier removal i.e., by rejecting points beyond a distance of 10 pixels from the mean of the boundary point locations. A polygon is constructed using the remainder of the point points which serves as the delineated boundary. With sufficiently densely sampled points, the nearest points are connected without and topological issues found in practice. The entire training algorithm is summarised in Algo. 1, where at inference, the only difference is that rewards are not computed and the controller is not updated.

3 Experiments

3.1 Dataset

During the (NCT02290561, NCT02341677) clinical trials, TRUS images from 249 patients were collected using a bi-plane transperineal ultrasound probe during manual/rotational positioning of a digital transperineal stepper, resulting in a total of 5185 prostate-containing images with the size of 403×361 , which were center-cropped to 360×360 . Ground truth labels were acquired through pixel-wise consensus among three researchers and verified by an expert radiologist. From 203 patients, we randomly selected 4000 images for controller training, and further cropped 80000 patches sized 90×90 for MIL and Boundary Presence Classifier training. From the remaining 46 patients, we randomly chose 200 images each for validation and testing.

3.2 Architecture and implementation details

Boundary presence classifier: EfficientNet [28] architecture was used, with the Adam optimiser with a learning rate of $1e-4$ and batch size 16.

Boundary RL controller: We trained the controller, which has 3 convolutional layers followed by 2 dense layers, for Boundary-RL using Proximal Policy Optimization (PPO) [25], which took about 27 hours on a single Nvidia Tesla V100 GPU, with a learning rate set at $3e-4$ and a batch size of 4096.

Data: Images to construct $s_t \in \mathcal{S}$
Result: Trained RL policy π_{θ^*} .

```

while not converged do
  Randomly sample an image  $x_t \in \mathcal{X}$ ;
  for  $m \leftarrow 1$  to  $M$  do
    Start at  $t = 0$ ;
    Randomly sample patch coordinates  $c_0 = (i, j)$ ;
    Construct the state  $s_t = \{x_t, x_{c,0}\}$ ;
    Sample the action  $a_0$  according to the policy  $a_0 \sim \pi_{\theta}(a_0|s_0)$ ;
    Compute object presence-based reward  $R_0 = r(s_0, a_0)$ ;

    for  $t \leftarrow 1$  to  $T$  do
      Note:  $t$  is now iterating starting at  $t = 1$ ;
      Given  $a_{t-1} = (\delta i, \delta j)$ , update patch coordinates  $c_t = (i + \delta i, j + \delta j)$ ;
      Construct the state  $s_t = \{x_t, x_{c,t}\}$ ;
      Sample the action  $a_t$  according to the policy  $a_t \sim \pi_{\theta}(a_t|s_t)$ ;
      Compute object presence-based reward  $R_t = r(s_t, a_t)$ ;
      End if object presence detected i.e.,  $t = t_{term}$  if  $f(x_{c,t}; w^*) = 1$ ;
      add Gaussian noise to patch to update  $x_t$ ;
    end
  end
  Once  $M$  trajectories of  $R_{t=1:T}$  or  $R_{t=1:t_{term}}$  collected, update RL function
  using gradient ascent
end

```

Algorithm 1: Training procedure for Boundary RL.

3.3 Comparisons and baselines

MIL: MIL is adapted from previous works [16, 17, 27, 11, 19], for our application of prostate gland segmentation. We treat each patch as a bag, and each pixel within a patch is considered an instance. We use a modified U-Net [20], with global average pooling added before the full connection layer to aggregate instance probability distribution vectors into bag features. We adopt both partition (MIL-P) and sliding window (MIL-S) approaches for prediction. For overlapping predictions within the sliding window, we take their average.


Fully supervised learning (FSL): We chose the commonly used U-Net, trained using segmentation labels, to compare with our proposed WSSS method.

4 Results

Results presented in Tab. 1, show that our proposed Boundary-RL attained highest performance amongst the tested WSSS methods. While performance was statistically significantly better than the commonly used MIL with image partitioning (MIL-P) for inference, significance was not observed for the comparison with MIL using sliding windows (MIL-S) at inference. The fully supervised (FSL) method performed better than all tested WSSS methods, with statistical significance. Samples from each of the methods are presented in Fig. 2.

Table 1. Table of results and corresponding statistical tests.

Algorithm	Dice	Statistical tests (t-value/ p-value)			
		Proposed	MIL - P	MIL - S	FSL
Proposed	0.751 ± 0.161	N/A	-	-	-
MIL - P	0.663 ± 0.156	5.53/ <0.01	N/A	-	-
MIL - S	0.737 ± 0.152	0.88/ 0.38	4.80/ <0.01	N/A	-
FSL	0.846 ± 0.167	5.76/ <0.01	11.26/ <0.01	6.79/ <0.01	N/A

**Fig. 2.** Samples with segmentation from the tested methods.

5 Discussion and Conclusion

Results show that an effective boundary delineation policy may be learnt using our proposed Boundary-RL training scheme. Proposed adaptations to the commonly used MIL approach, also yielded comparable performance, with Boundary-RL being potentially more time-efficient at inference due to a fewer passes through the classifier due to the controller-learnt localisation policy, compared to sliding windows in MIL which may be akin to brute-search. All WSSS methods under-performed fully-supervised learning as expected due to the use of pixel-level segmentation maps for training, compared to only binary classification labels used in WSSS. Boundary-RL outperformed all tested WSSS methods but statistical significance was not observed compared to MIL-S at inference. We maintain that our method produces an effective boundary delineation, evidenced from qualitative results and high observed Dice. Qualitative samples show that Boundary-RL can handle cases where artefacts exist within ROI boundaries, where MIL was unable to effectively segment since it relies on pixel-level classification which may be impacted by presence of artefacts within ROI boundaries. Future studies could explore augmenting rewards with other weak supervision signals such as self re-construction instead of just the patch-level classification.

We proposed a WSSS method based on RL that envisions segmentation as a boundary detection problem. A MDP is formulated to effectively learn to localise ROI boundaries using a reward based on boundary-presence within the localised region. The reward, generated using a classifier trained using binary patch-level labels of object presence, serves as a weak signal to inform boundary detection. We evaluate our approach for a clinically relevant task of boundary delineation of the prostate gland on TRUS, used during navigation in TRUS guided biopsy.

Acknowledgements: This work was supported by the EPSRC grant [EP/T029404/1], Wellcome/EPSRC Centre for Interventional and Surgical Sciences [203145Z/16/Z], and the International Alliance for Cancer Early Detection, an alliance between Cancer Research UK [C28070/A30912; 73666/A31378], Canary Center at Stanford University, the University of Cambridge, OHSU Knight Cancer Institute, University College London and the University of Manchester.

References

1. Chalcraft, L.F., Qu, J., Martin, S.A., Gayo, I.J., et al.: Development and evaluation of intraoperative ultrasound segmentation with negative image frames and multiple observer labels. In: ASMUS 2021. Springer (2021)
2. Chen, M.Y., Woodruff, M.A., Dasgupta, P., Rukin, N.J.: Variability in accuracy of prostate cancer segmentation among radiologists, urologists, and scientists. *Cancer Medicine* **9**(19) (2020)
3. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation policies from data. arXiv preprint arXiv:1805.09501 (2018)
4. Czolbe, S., Arnavaz, K., Krause, O., Feragen, A.: Is segmentation uncertainty useful? In: IPMI 2021. Springer (2021)
5. Dertkigil, S., Appenzeller, S., Lotufo, R., Rittner, L.: A systematic review of automated segmentation methods and public datasets for the lung and its lobes and findings on computed tomography images. *Yearbook of Medical Informatics* **31**(01) (2022)
6. Han, C., Lin, J., Mai, J., Wang, Y., Zhang, Q., et al.: Multi-layer pseudo-supervision for histopathology tissue semantic segmentation using patch-level classification labels. *MedIA* **80** (2022)
7. Hulsmans, F.j.J., Castelijns, J.A., Reeders, J.W., Tytgat, G.N.: Review of artifacts associated with transrectal ultrasound: understanding, recognition, and prevention of misinterpretation. *Journal of clinical ultrasound* **23**(8) (1995)
8. Jia, Z., Huang, X., Eric, I., Chang, C., Xu, Y.: Constrained deep weak supervision for histopathology image segmentation. *IEEE TMI* **36**(11) (2017)
9. Karimi, D., Zeng, Q., Mathur, P., Avinash, A., et al.: Accurate and robust deep learning-based segmentation of the prostate clinical target volume in ultrasound images. *MedIA* **57** (2019)
10. Kervadec, H., Dolz, J., Tang, M., Granger, E., Boykov, Y., et al.: Constrained-cnn losses for weakly supervised segmentation. *MedIA* **54** (2019)
11. Kots, M., Chukanov, V.: U-net adaptation for multiple instance learning. In: *Journal of Physics: Conference Series*. vol. 1236. IOP Publishing (2019)
12. Lei, Y., Tian, S., He, X., Wang, T., et al.: Ultrasound prostate segmentation based on multidirectional deeply supervised v-net. *Medical physics* **46**(7) (2019)
13. Li, K., Qian, Z., Han, Y., Eric, I., Chang, C., et al.: Weakly supervised histopathology image segmentation with self-attention. *MedIA* **86** (2023)
14. Liu, D., Wang, L., Du, Y., Cong, M., Li, Y.: 3-d prostate mr and trus images detection and segmentation for puncture biopsy. *IEEE Transactions on Instrumentation and Measurement* **71** (2022)
15. Pathak, D., Krahenbuhl, P., Darrell, T.: Constrained convolutional neural networks for weakly supervised segmentation. In: *ICCV* (2015)
16. Pathak, D., Shelhamer, E., Long, J., Darrell, T.: Fully convolutional multi-class multiple instance learning. arXiv preprint arXiv:1412.7144 (2014)

17. Pinheiro, P.O., Collobert, R.: From image-level to pixel-level labeling with convolutional networks. In: CVPR (2015)
18. Ramesh, K., Kumar, G.K., Swapna, K., Datta, D., Rajest, S.S.: A review of medical image segmentation algorithms. *EAI PHAT* **7**(27) (2021)
19. Ren, W., Huang, K., Tao, D., Tan, T.: Weakly supervised large scale object localization with multiple instance learning and bag splitting. *IEEE transactions on pattern analysis and machine intelligence* **38**(2) (2015)
20. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI 2015. Springer (2015)
21. Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M., Yang, Q., et al.: Adaptable image quality assessment using meta-reinforcement learning of task amenability. In: ASMUS 2021. Springer (2021)
22. Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M., Yang, Q., et al.: Image quality assessment for machine learning tasks using meta-reinforcement learning. *MedIA* **78** (2022)
23. Saeed, S.U., Yan, W., Fu, Y., Giganti, F., et al.: Image quality assessment by overlapping task-specific and task-agnostic measures: application to prostate multiparametric MR images for cancer segmentation. *Machine Learning for Biomedical Imaging 1(IPMI 2021)* (2022)
24. Savjani, R.R., Lauria, M., Bose, S., Deng, J., et al.: Automated tumor segmentation in radiotherapy. In: *Seminars in radiation oncology*. vol. 32. Elsevier (2022)
25. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017)
26. Sharma, N., Aggarwal, L.M.: Automated medical image segmentation techniques. *Journal of medical physics* **35**(1) (2010)
27. Shi, X., Xing, F., Xie, Y., Zhang, Z., Cui, L., et al.: Loss-based attention for deep multiple instance learning. In: *AAAI 2020*. vol. 34 (2020)
28. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International conference on machine learning*. PMLR (2019)
29. Wang, X., Chang, Z., Zhang, Q., Li, C., Miao, F., et al.: Prostate ultrasound image segmentation based on dsu-net. *Biomedicines* **11**(3) (2023)
30. Wang, Y., Dou, H., Hu, X., Zhu, L., Yang, X., et al.: Deep attentive features for prostate segmentation in 3d transrectal ultrasound. *IEEE TMI* **38**(12) (2019)
31. Xu, G., Song, Z., Sun, Z., Ku, C., et al.: Camel: A weakly supervised learning framework for histopathology image segmentation. In: *IEEE/ CVF ICCV* (2019)
32. Xu, R.S., Michailovich, O., Salama, M.: Information tracking approach to segmentation of ultrasound imagery of the prostate. *IEEE transactions on ultrasonics, ferroelectrics, and frequency control* **57**(8) (2010)
33. Xu, Y., Zhu, J.Y., Eric, I., Chang, C., Lai, M., et al.: Weakly supervised histopathology cancer image segmentation and classification. *MedIA* **18**(3) (2014)
34. Yoon, J., Arik, S., Pfister, T.: Data valuation using reinforcement learning. In: *International Conference on Machine Learning*. PMLR (2020)
35. Zeng, Y., Tsui, P.H., Wu, W., Zhou, Z., Wu, S.: Fetal ultrasound image segmentation for automatic head circumference biometry using deeply supervised attention-gated v-net. *Journal of Digital Imaging* **34** (2021)
36. Zhang, M., Zhou, Y., Zhao, J., Man, Y., Liu, B., et al.: A survey of semi-and weakly supervised semantic segmentation of images. *AIRE* **53** (2020)
37. Zhang, X., Wang, Q., Zhang, J., Zhong, Z.: Adversarial autoaugment. *arXiv preprint arXiv:1912.11188* (2019)