

DeepPyramid: Enabling Pyramid View and Deformable Pyramid Reception for Semantic Segmentation in Cataract Surgery Videos^{*}

Negin Ghamsarian¹, Mario Taschwer², Raphael Sznitman¹, and Klaus Schoeffmann²

¹ Center for AI in Medicine, Faculty of Medicine, University of Bern
{negin.ghamsarian, raphael.sznitman}@unibe.ch

² Department of Information Technology, Alpen-Adria-Universität Klagenfurt
{mt,ks}@itec.aau.at

Abstract. Semantic segmentation in cataract surgery has a wide range of applications contributing to surgical outcome enhancement and clinical risk reduction. However, the varying issues in segmenting the different relevant structures in these surgeries make the designation of a unique network quite challenging. This paper proposes a semantic segmentation network, termed DeepPyramid, that can deal with these challenges using three novelties: (1) a Pyramid View Fusion module which provides a varying-angle global view of the surrounding region centering at each pixel position in the input convolutional feature map; (2) a Deformable Pyramid Reception module which enables a wide deformable receptive field that can adapt to geometric transformations in the object of interest; and (3) a dedicated Pyramid Loss that adaptively supervises multi-scale semantic feature maps. Combined, we show that these modules can effectively boost semantic segmentation performance, especially in the case of transparency, deformability, scalability, and blunt edges in objects. We demonstrate that our approach performs at a state-of-the-art level and outperforms a number of existing methods with a large margin (3.66% overall improvement in intersection over union compared to the best rival approach).

Keywords: Cataract Surgery· Semantic Segmentation· Surgical Data Science

1 Introduction

Cataracts are naturally developing opacity that obfuscates sight and is the leading cause of blindness worldwide, with over 100 million people suffering from them. Today, surgery is the most effective way to cure patients by replacing natural eye lenses with artificial ones. More than 10 million cataract surgeries are performed every year, making it one of the most common surgeries globally [20]. With the aging world population growing, the number of patients at risk of complete cataract-caused blindness is sharply increasing [21] and the number of surgeries needed brings unprecedented organizational and logistical challenges.

^{*} This work was funded by Haag-Streit Switzerland and the FWF Austrian Science Fund under grant P 31486-N31.

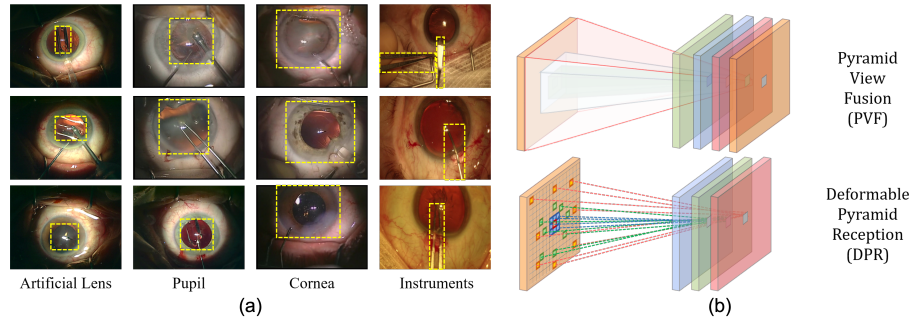


Fig. 1: (a) Challenges in semantic segmentation of relevant objects in cataract surgery. (b) Proposed Pyramid View Fusion and Deformable Pyramid Reception modules.

To help train future surgeons and optimize surgical workflows, automated methods that analyze cataract surgery videos have gained significant traction in the last decade. With the prospect of reducing intra-operative and post-operative complications [5], recent methods have included surgical skill assessment [8,26], remaining surgical time estimation [13], irregularity detection [7] or relevance-based compression [6]. In addition, a reliable relevant-instance-segmentation approach is often a prerequisite for a majority of these applications [17]. In this regard, four different structures are typically of interest: the intraocular lens, the pupil, the cornea, and surgical instruments. Due to the diversity in the appearance of these structures, segmentation methods must overcome several hurdles to perform well on real-world video sequences. Specifically, a semantic segmentation network is required to simultaneously deal with: 1) a transparent artificial lens that undergoes deformations, 2) color, shape, size, and texture variations in the pupil, 3) unclear edges of the cornea, and 4) severe motion blur, reflection distortion, and scale variations in instruments (see Fig. 1-a). This work looks to provide a method to segment these structures despite the mentioned challenges.

Several network architectures for cataract surgery semantic segmentation have been proposed or have been used in the recent past [16,14,15,1,24]. Many of these methods have been based on the U-Net architecture [19] and aimed at improving accuracy by addressing different limitations from the original architecture. In [16,14], different attention modules were used to guide the network’s computational efforts toward the most discriminative features in the input feature map considering the characteristics of the objects of interest. Additionally, fusion modules have been proposed to improve semantic representation via combining several feature maps [1,24]. However, as we show in our experiments, these methods still have difficulties with the aforementioned challenges in cataract video segmentation.

In this work, we propose a novel architecture that is tailored to adaptively capture semantic information despite the challenges typically found in cataract surgery videos. Our proposed network, *DeepPyramid*¹, introduces three key contributions: (i) a Pyra-

¹ The PyTorch implementation of DeepPyramid is publicly available at https://github.com/Negin-Ghamsarian/DeepPyramid_MICCAI2022




Fig. 2: Overall architecture of the proposed DeepPyramid network. It contains Pyramid View Fusion (PVF), Deformable Pyramid Reception (DPR), and Pyramid Loss (PL).

mid View Fusion (PVF) module allowing a varying-angle surrounding view of the feature maps for each pixel position, (ii) a Deformable Pyramid Reception (DPR) module, which enables a large, sparse, and learnable receptive field to perform shape-wise feature extraction (see Fig. 1-b), and (iii) a Pyramid Loss, (PL) to explicitly supervise multi-scale semantic feature maps in our network. We show in the experiments that our approach outperforms by a significant margin twelve rival state-of-the-art approaches for cataract surgery segmentation. Specifically, we show that our model is particularly effective for deformable, transparent, and changing scale objects. In addition, we show the contribution of each of the proposed additions, highlighting that the addition of all three yields the observed improvements.

2 Methodology

Our proposed segmentation strategy aims to explicitly model deformations and context within its architecture. Using a U-Net-based architecture, our proposed model is illustrated in Fig. 2. At its core, the encoder network remains that of a standard VGG16 network. Our approach is to provide useful decoder modules to help alleviate segmentation concerning relevant objects' features in cataract surgery². Specifically, we propose a Pyramid View Fusion (PVF) module and a Deformable Pyramid Reception (DPR) module. These are then trained using a dedicated Pyramid Loss (PL).

Conceptually, the PVF module is inspired by the human visual system and aims to recognize semantic information found in images considering not only the internal object's content but also the relative information between the object and its surrounding area. Thus the role of the PVF is to reinforce the observation of relative information at every distinct pixel position. Specifically, we use average pooling to fuse the multi-angle local information for this novel attention mechanism. Conversely, our DPR module hinges on a novel deformable block based on dilated convolutions that can help recognize each pixel position's semantic label based on its cross-dependencies with varying-distance surrounding pixels without imposing additional trainable parameters.

² Since changing the encoder network entails pretraining on a large dataset (such as ImageNet), which in turn imposes more computational costs, we only add the proposed modules after the bottleneck. Nevertheless, since these modules are applied to concatenated features coming from the encoder network via skip connections, the encoder features can be effectively guided.

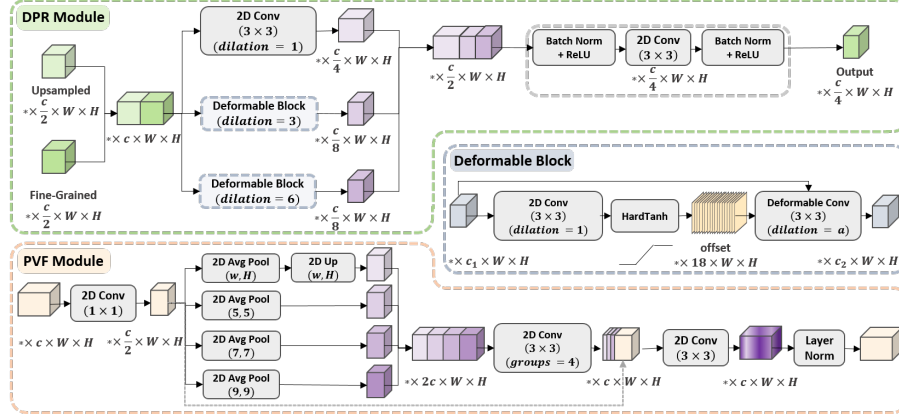


Fig. 3: Detailed architecture of the Deformable Pyramid Reception (DPR) and Pyramid View Fusion (PVF) modules.

Due to the inflexible rectangle shape of the receptive field in regular convolutional layers, the feature extraction procedure cannot be adapted to complex deformable shapes [12]. Our proposed dilated deformable convolutional layers attempt to remedy this explicitly in terms of both scale and shape. We now specify these modules and our loss function in the following subsections.

Pyramid View Fusion (PVF). First, a bottleneck is formed by employing a convolutional layer with a kernel size of one to curb computational complexity. The convolutional feature map is then fed into four parallel branches: a global average pooling layer followed by upsampling and three average pooling layers with progressively larger filter sizes and a common stride of 1. Note that using a one-pixel stride is essential to obtain pixel-wise centralized pyramid views in contrast with region-wise pyramid attention as shown in PSPNet [24]. The output feature maps are then concatenated and fed into a convolutional layer with four groups. This layer is responsible for extracting inter-channel dependencies during dimensionality reduction. A regular convolutional layer is then applied to extract joint intra-channel and inter-channel dependencies before being fed into a layer-normalization function. A summary of this module is illustrated in Fig. 3.

Deformable Pyramid Reception (DPR). As shown in Fig. 3 (top), the fine-grained feature map from the encoder and coarse-grained semantic feature map from the previous layer are first concatenated. These features are then fed into three parallel branches: a regular 3×3 convolution and two deformable blocks with different dilation rates. Together, these layers cover a learnable but sparse receptive field of size 15×15 ³ as

³ The structured 3×3 filter covers up to 1 pixel from the central pixel. The deformable filter with $dilation = 3$ covers an area of 2 to 4 pixels away from each central pixel. Similarly, the deformable convolution with $dilation = 6$ covers an area of 5 to 7 pixels away from each central pixel. Together, these form a sparse filter of size 15×15 pixels.

shown in Fig. 1 (b). The output feature maps are then concatenated before undergoing a sequence of regular layers for higher-order feature extraction and dimensionality reduction.

The deformable blocks used in the DPR module consist of a regular convolutional layer applied to the input feature map to compute an offset field for deformable convolution. The offset field provides two values per element in the convolutional filter (horizontal and vertical offsets). Accordingly, the number of offset field’s output channels for a kernel of size 3×3 is equal to 18. Inspired by dU-Net [23], the convolutional layer for the offset field is followed by an activation function, which we set to the hard tangent hyperbolic function, as it is computationally efficient and clips offset values to the range of $[-1, 1]$. In summary (see Fig. 3), the deformable block uses learned offset values along with the convolutional feature map with a predetermined dilation rate to extract object-adaptive features.

The output feature map (y) for each pixel position (p_0) and the receptive field (\mathcal{RF}) for a regular 2D convolution with a 3×3 filter and dilation rate of 1 can be computed by,

$$y(p_0) = \sum_{p_i \in \mathcal{RF}_1} x(p_0 + p_i)w(p_i), \quad (1)$$

where $\mathcal{RF}_1 = \{(-1, -1), (-1, 0), \dots, (1, 0), (1, 1)\}$, x denotes the input convolutional feature map, and w refers to the weights of the convolutional kernel. In a dilated 2D convolution with a dilation rate of α , the receptive field can be defined as $\mathcal{RF}_\alpha = \alpha \times \mathcal{RF}_1$. Although the sampling locations in a dilated receptive field have a greater distance to the central pixel, they follow a firm structure. In a deformable dilated convolution with a dilation rate of α , the sampling locations of the receptive field are dependent on the local contextual features. In the proposed deformable block, the sampling location for the i th element of the receptive field and the input pixel p_0 are calculated as,

$$\mathcal{RF}_{def,\alpha}[i, p_0] = \mathcal{RF}_\alpha[i] + f \left(\sum_{p_j \in \mathcal{RF}_1} x(p_0 + p_j)\hat{w}(p_j) \right), \quad (2)$$

where f denotes the activation function, which is the tangent hyperbolic function in our case, and \hat{w} refers to the weights of the offset filter. This learnable receptive field can be adapted to every distinct pixel in the convolutional feature map and allows the convolutional layer to extract stronger informative semantic features when compared to the regular convolution.

Pyramid Loss (PL). To train our network using the PVF and DPR modules, we wish to directly supervise the multi-scale semantic feature maps of the decoder. To enable direct supervision, a depth-wise fully connected layer is formed using a pixel-wise convolution operation. The output feature map presents the semantic segmentation results with the same resolution as the input feature map. To compute the loss for varying-scale outputs, we downscale the ground-truth masks using inter-nearest downsampling for multi-class segmentation and max-pooling for binary segmentation. Our overall loss is then defined as,

$$PL = \mathcal{L}_1 + \alpha\mathcal{L}_2 + \beta\mathcal{L}_4 + \gamma\mathcal{L}_8, \quad (3)$$

where α , β , and γ are predetermined weights in the range of $[0, 1]$ and \mathcal{L}_i denotes the loss of output mask segmentation result with the resolution of $(1/i)$ compared to the input resolution.

3 Experimental Setup

To evaluate the performance of our approach, we make use of data from three datasets. These include the “Cornea” [6] and “Instruments” mask annotations from the CaDIS dataset [9]. In addition, we have collected a separate dataset from which we performed the “Intraocular Lens” and “Pupil” pixel-wise segmentations⁴. The total number of training and test images for the aforementioned objects are 178:84, 3190:459, 141:48, and 141:48, respectively⁵. In the following experiments, all training and test images were split patient-wise to ensure realistic conditions.

We compare the performance of DeepPyramid with thirteen different state-of-the-art segmentation approaches including UNet++ and UNet++/DS [25], CPFNet [4], BARNet [15], PAANet [14], dU-Net⁶ [23], MultiResUNet [11], CE-Net [10], RAUNet [16], FED-Net [2] UPerNet [22], PSPNet⁷ [24], and U-Net [19]⁸. With the exception of the U-Net, MultiResUNet, and dU-Net, which do not use a pretrained backbone, the weights of the backbone for all networks were initialized with ImageNet [3] training weights. The input size of all models is set to $3 \times 512 \times 512$.

For all methods, training is performed using data augmentation. Transformations considered the inherent and statistical features of datasets. For instance, we use motion blur transformation to encourage the network to deal with harsh motion blur regularly occurring in cataract surgery videos. We further use brightness and contrast, shift and scale, and rotate augmentation.

Due to the different depth and connections of the proposed and rival approaches, all networks are trained with three different initial learning rates ($lr \in \{5, 2, 10\} \times 10^{-4}$), and the results with the highest IoU for each network are listed. The learning rate is scheduled to decrease every two epochs with the factor of 0.8. In all evaluations, the networks are trained end-to-end and for 30 epochs. We use a threshold of 0.1 for gradient clipping during training.

The loss function used during training is a weighted sum of binary cross-entropy (*BCE*) and the logarithm of the soft Dice coefficient. We set $\alpha = 0.75$, $\beta = 0.5$, and $\gamma = 0.25$ in equation (3). Additional information on our experimental section can be found in the supplementary materials.

⁴ The customized datasets is publicly released in <https://ftp.itec.aau.at/datasets/ovid/DeepPyramid/>.

⁵ Our evaluations are based on binary segmentation per relevant object so that we do not have the imbalance problem. In the case of multi-class classification, methods such as oversampling can mitigate the imbalance problem [18].

⁶ Our version of du-Net has the same number of filter-response maps as the U-Net.

⁷ To provide a fair comparison, we adopt our improved version of PSPNet, featuring a decoder designed similarly to U-Net (with four sequences of double-convolution blocks).

⁸ BARNet, PAANet, and RAUNet are tailored for instrument segmentation in surgical videos. Other methods are state-of-the-art for medical image segmentation.

Table 1: Quantitative comparison between the proposed DeepPyramid and state-of-the-art approaches.

Network	IoU% Dice%									
	Lens		Pupil		Cornea		Instrument		Mean	
U-Net [19]	58.19	67.91	85.51	89.36	79.83	86.20	56.12	67.02	69.91	77.62
PSPNet+ [24]	80.56	88.89	93.23	96.45	88.09	93.55	65.37	76.47	81.81	88.84
UPerNet [22]	77.78	86.93	93.34	96.52	86.62	92.67	68.51	78.68	81.56	88.70
FEDNet [2]	78.12	87.38	93.93	96.85	85.73	92.10	65.13	76.11	80.72	88.11
RAUNet [16]	76.40	85.34	89.26	94.26	85.73	92.10	65.13	76.11	79.13	86.95
CE-Net [10]	68.40	80.43	83.59	90.89	83.47	90.85	61.57	74.64	74.25	84.20
MultiResUNet [11]	60.73	71.62	58.36	66.80	73.10	83.40	55.43	66.07	61.90	71.97
dU-Net [23]	59.83	69.46	71.86	79.53	82.39	90.00	61.36	71.55	68.86	77.63
PAANet [14]	74.92	84.83	90.02	94.59	86.75	92.71	64.47	75.24	79.04	86.74
BARNet [15]	67.33	78.85	91.33	95.32	83.98	91.09	66.72	77.14	77.34	85.60
CPFNet [4]	73.56	83.74	90.27	94.83	87.63	93.28	61.16	73.51	78.18	86.34
UNet++/DS [25]	79.50	87.85	95.28	97.53	86.72	92.57	66.05	75.91	81.88	88.46
UNet++ [25]	81.32	89.34	95.66	97.77	85.08	91.72	70.11	79.56	83.04	89.59
DeepPyramid	85.61	91.98	96.56	98.24	90.24	94.63	74.40	83.30	86.70	92.03

4 Experimental Results

Table 1 compares the performance of all evaluated methods. Accordingly, DeepPyramid, Unet++, and PSPNet+ are the top three segmentation methods in terms of IoU for the relevant objects in cataract surgery videos. However, DeepPyramid shows considerable improvements in segmentation accuracy compared to the second-best approach in each class. Specifically, DeepPyramid achieves more than 4% improvement in lens segmentation (85.61% vs. 81.32%) and more than 4% improvement in instrument segmentation (74.40% vs. 70.11%) compared to UNet++. Similarly, DeepPyramid achieves the highest dice coefficient compared to the evaluated approaches for all classes.

Table 2 validates the effectiveness of the proposed modules in an ablation study while also showing the impact on the different segmentation classes. The PVF module appears to enhance the performance for the cornea and instrument segmentation (2.41% and 2.76% improvement in IoU, respectively). This improvement is most likely due to the ability of the PVF module to provide a global view of varying-size sub-regions centered around each spatial position. Such a global view can reinforce semantic representation in the regions corresponding to blunt edges and reflections. Due to scale variance in instruments, the DPR module boosts the segmentation performance for surgical instruments. The addition of the $P\mathcal{L}$ loss results in the improvement in IoU for all the relevant classes, especially the lens (roughly 2% improvement) and instrument (1.64% improvement) classes. The combination of PVF, DPR, and $P\mathcal{L}$ show a marked 4.58% improvement in instrument segmentation and 4.22% improvement in cornea segmentation (based on IoU%). These modules improve the IoU for the lens and pupil by 2.85% and 1.43%, respectively. Overall, the addition of the different proposed components in DeepPyramid lead to considerable improvements in segmentation performance (3.27% improvement in IoU) when compared to the evaluated baselines.

Table 2: Impact of different modules on DeepPyramid’s performance (ablation study).

Modules				IoU%/Dice%				
PVF	DPR	$P\mathcal{L}$	Params	Lens	Pupil	Cornea	Instrument	Overall
\times	\times	\times	22.55 M	82.98/ 90.44	95.13/ 97.48	86.02/ 92.28	69.82/ 79.05	83.49/ 89.81
\checkmark	\times	\times	22.99 M	83.73/ 90.79	96.04/ 97.95	88.43/ 93.77	72.58/ 81.84	85.19/ 91.09
\times	\checkmark	\times	23.17 M	81.85/ 89.58	95.32/ 97.59	86.43/ 92.55	71.57/ 80.60	83.79/ 90.08
\checkmark	\checkmark	\times	23.62 M	83.85/ 90.89	95.70/ 97.79	89.36/ 94.29	72.76/ 82.00	85.42/ 91.24
\checkmark	\checkmark	\checkmark	23.62 M	85.84/ 91.98	96.56/ 98.24	90.24/ 94.77	74.40/ 83.30	86.76/ 92.07




Fig. 4: Qualitative comparisons among DeepPyramid and the six top-performing rival approaches for the relevant objects in cataract surgery videos. The representative images are selected from the test set.

Fig. 4 illustrates the qualitative results of our method and evaluated baselines. Specifically, we see the effectiveness DeepPyramid has in segmenting challenging cases. Taking advantage of the pyramid view provided by the PVF module, DeepPyramid can handle reflection and brightness variation in instruments, blunt edges in the cornea, color and texture variation in the pupil, as well as transparency in the lens. Furthermore, powered by deformable pyramid reception, DeepPyramid can tackle scale variations in instruments and blunt edges in the cornea. In particular, we see from Fig. 4 that DeepPyramid shows much less distortion in the region of edges, especially in the case of the cornea. Furthermore, based on these qualitative experiments, DeepPyramid shows much better precision and recall in the narrow regions for segmenting the instruments and other relevant objects in the case of occlusion by the instruments. Further results are shown in the supplementary materials of the paper.

5 Conclusion

In this work, we have proposed a novel network architecture for semantic segmentation in cataract surgery videos. The proposed architecture takes advantage of two modules, namely “Pyramid View Fusion” and “Deformable Pyramid Reception”, as well as a dedicated “Pyramid Loss”, to simultaneously deal with (i) geometric transformations such as scale variation and deformability, (ii) blur degradation and blunt edges, and (iii) transparency, texture and color variation typically observed in cataract surgery images. We show in our experiments that our approach provides state-of-the-art performances in segmenting key anatomical structures and surgical instruments typical with such surgeries. Beyond this, we demonstrate that our approach outperforms a large number of recent segmentation methods by a considerable margin. The proposed architecture can also be adopted for various other medical image segmentation and general semantic segmentation problems.

References

1. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European Conference on Computer Vision (ECCV) (September 2018)
2. Chen, X., Zhang, R., Yan, P.: Feature fusion encoder decoder network for automatic liver lesion segmentation. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 430–433 (2019). <https://doi.org/10.1109/ISBI.2019.8759555>
3. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
4. Feng, S., Zhao, H., Shi, F., Cheng, X., Wang, M., Ma, Y., Xiang, D., Zhu, W., Chen, X.: Cpfnet: Context pyramid fusion network for medical image segmentation. IEEE Transactions on Medical Imaging **39**(10), 3008–3018 (2020). <https://doi.org/10.1109/TMI.2020.2983721>
5. Ghamsarian, N.: Enabling relevance-based exploration of cataract videos. In: Proceedings of the 2020 International Conference on Multimedia Retrieval. p. 378–382. ICMR ’20 (2020). <https://doi.org/10.1145/3372278.3391937>
6. Ghamsarian, N., Amirpourazarian, H., Timmerer, C., Taschwer, M., Schöffmann, K.: Relevance-based compression of cataract surgery videos using convolutional neural networks. In: Proceedings of the 28th ACM International Conference on Multimedia. p. 3577–3585. MM ’20, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3394171.3413658>, <https://doi.org/10.1145/3394171.3413658>
7. Ghamsarian, N., Taschwer, M., Putzgruber-Adamitsch, D., Sarny, S., El-Shabrawi, Y., Schoeffmann, K.: Lensid: A cnn-rnn-based framework towards lens irregularity detection in cataract surgery videos. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. pp. 76–86. Springer International Publishing, Cham (2021)
8. Ghamsarian, N., Taschwer, M., Putzgruber-Adamitsch, D., Sarny, S., Schoeffmann, K.: Relevance detection in cataract surgery videos by spatio-temporal action localization (2021)
9. Grammatikopoulou, M., Flouty, E., Kadkhodamohammadi, A., Quelled, G., Chow, A., Nehme, J., Luengo, I., Stoyanov, D.: Cadis: Cataract dataset for image segmentation (2020)

10. Gu, Z., Cheng, J., Fu, H., Zhou, K., Hao, H., Zhao, Y., Zhang, T., Gao, S., Liu, J.: Ce-net: Context encoder network for 2d medical image segmentation. *IEEE Transactions on Medical Imaging* **38**(10), 2281–2292 (2019). <https://doi.org/10.1109/TMI.2019.2903562>
11. Ibtehaz, N., Rahman, M.S.: Multiresunet : Rethinking the u-net architecture for multimodal biomedical image segmentation. *Neural Networks* **121**, 74–87 (2020). <https://doi.org/https://doi.org/10.1016/j.neunet.2019.08.025>, <https://www.sciencedirect.com/science/article/pii/S0893608019302503>
12. Lei, T., Wang, R., Zhang, Y., Wan, Y., Liu, C., Nandi, A.K.: Defed-net: Deformable encoder-decoder network for liver and liver tumor segmentation. *IEEE Transactions on Radiation and Plasma Medical Sciences* **6**(1), 68–78 (2022). <https://doi.org/10.1109/TRPMS.2021.3059780>
13. Marafioti, A., Hayoz, M., Gallardo, M., Márquez Neila, P., Wolf, S., Zinkernagel, M., Sznitman, R.: Catanet: Predicting remaining cataract surgery duration. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 426–435 (2021)
14. Ni, Z.L., Bian, G.B., Wang, G.A., Zhou, X.H., Hou, Z.G., Chen, H.B., Xie, X.L.: Pyramid attention aggregation network for semantic segmentation of surgical instruments. *Proceedings of the AAAI Conference on Artificial Intelligence* **34**(07), 11782–11790 (Apr 2020). <https://doi.org/10.1609/aaai.v34i07.6850>, <https://ojs.aaai.org/index.php/AAAI/article/view/6850>
15. Ni, Z.L., Bian, G.B., Wang, G.A., Zhou, X.H., Hou, Z.G., Xie, X.L., Li, Z., Wang, Y.H.: Barnet: Bilinear attention network with adaptive receptive fields for surgical instrument segmentation. In: Bessiere, C. (ed.) *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*. pp. 832–838 (7 2020). <https://doi.org/10.24963/ijcai.2020/116>
16. Ni, Z.L., Bian, G.B., Zhou, X.H., Hou, Z.G., Xie, X.L., Wang, C., Zhou, Y.J., Li, R.Q., Li, Z.: Raunet: Residual attention u-net for semantic segmentation of cataract surgical instruments. In: Gedeon, T., Wong, K.W., Lee, M. (eds.) *Neural Information Processing* (2019)
17. Pissas, T., Ravasio, C.S., Da Cruz, L., Bergeles, C.: Effective semantic segmentation in cataract surgery: What matters most? In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 509–518. Springer International Publishing, Cham (2021)
18. Pissas, T., Ravasio, C.S., Da Cruz, L., Bergeles, C.: Effective semantic segmentation in cataract surgery: What matters most? In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 509–518. Springer International Publishing, Cham (2021)
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. pp. 234–241 (2015)
20. Trikha, S., Turnbull, A., Morris, R., Anderson, D., Hossain, P.: *The journey to femtosecond laser-assisted cataract surgery: New beginnings or a false dawn?* *Eye* (London, England) **27** (02 2013). <https://doi.org/10.1038/eye.2012.293>
21. Wang, W., Yan, W., Müller, A., He, M.: A global view on output and outcomes of cataract surgery with national indices of socioeconomic development. *Investigative Ophthalmology and Visual Science* **58**, 3669–3676 (07 2017). <https://doi.org/10.1167/iovs.17-21489>
22. Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J.: Unified perceptual parsing for scene understanding. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *Computer Vision – ECCV 2018*. pp. 432–448. Springer International Publishing, Cham (2018)
23. Zhang, M., Li, X., Xu, M., Li, Q.: Automated semantic segmentation of red blood cells for sickle cell disease. *IEEE Journal of Biomedical and Health Informatics* **24**(11), 3095–3102 (2020). <https://doi.org/10.1109/JBHI.2020.3000484>
24. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (July 2017)

25. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: Unet++: Redesigning skip connections to exploit multiscale features in image segmentation. *IEEE Transactions on Medical Imaging* **39**(6), 1856–1867 (2020). <https://doi.org/10.1109/TMI.2019.2959609>
26. Zisimopoulos, O., Flouty, E., Luengo, I., Giataganas, P., Nehme, J., Chow, A., Stoyanov, D.: Deepphase: Surgical phase recognition in cataracts videos. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*. pp. 265–272. Springer International Publishing, Cham (2018)