Published in final edited form as:

Conf Proc IEEE Eng Med Biol Soc. 2019 July ; 2019: 7205–7211. doi:10.1109/EMBC.2019.8856791.

U-NetPlus: A Modified Encoder-Decoder U-Net Architecture for Semantic and Instance Segmentation of Surgical Instruments from Laparoscopic Images

S. M. Kamrul Hasan [Student Member, IEEE],

Center for Imaging Science, Rochester Institute of Technology, Rochester, NY USA.

Cristian A. Linte [Senior Member, IEEE]

Biomedical Engineering and Center for Imaging Science, Rochester Institute of Technology, Rochester, NY USA.

Abstract

With the advent of robot-assisted surgery, there has been a paradigm shift in medical technology for minimally invasive surgery. However, it is very challenging to track the position of the surgical instruments in a surgical scene, and accurate detection & identification of surgical tools is paramount. Deep learning-based semantic segmentation in frames of surgery videos has the potential to facilitate this task. In this work, we modify the U-Net architecture by introducing a pre-trained encoder and re-design the decoder part, by replacing the transposed convolution operation with an upsampling operation based on nearest-neighbor (NN) interpolation. To further improve performance, we also employ a very fast and flexible data augmentation technique. We trained the framework on 8×225 frame sequences of robotic surgical videos available through the MICCAI 2017 EndoVis Challenge dataset and tested it on 8×75 frame and 2×300 frame videos. Using our U-NetPlus architecture, we report a 90.20% DICE for binary segmentation, 76.26% DICE for instrument part segmentation, and 46.07% for instrument type (i.e., all instruments) segmentation, outperforming the results of previous techniques implemented and tested on these data.

I. Introduction

Minimally invasive surgery has addressed many of the challenges of traditional surgical approaches by significantly reducing the risk of infections and shortening hospitalization times, achieving a similar outcome to traditional open surgery. There is a new paradigm shift in this field thanks to robot assistance under laparoscopic visualization [19]. To facilitate the manipulation of the laparoscopic surgical instruments while visualizing the endoscopic scene, surgical instrument identification is critical. Nevertheless, this task is challenging, due to the surrounding effects like illumination changes, visual occlusions, and presence of nonclass objects. Hence, it is important to devise segmentation techniques that are sufficiently accurate and robust to ensure accurate tracking of the surgical tools to facilitate therapy via accurate manipulation of the laparoscopic instruments.

Although in recent years semantic segmentation methods applied to city-scapes, street scenes, and even Landsat image datasets [8], [25] have achieved ground-breaking performance by the virtue of deep convolutional neural networks (CNNs), image segmentation in clinical settings still requires more accuracy and precision, with even minimal segmentation errors being unacceptable.

The first fully convolutional network was proposed by Long *et al.* [14] for semantic segmentation. However, because of the limited size of the training dataset, its use in the medical domain has been challenging. Several techniques, including transfer learning [23], data augmentation and patch-based training [15], have been devised in the effort to mitigate the above challenge. However, semantic segmentation is not sufficiently accurate for handling multi-class objects, due to the close presence of objects of the same class in the surgical scene. Therefore, the proposed work is motivated by the need to improve multi-class object segmentation, by leveraging the power of the existing U-Net architecture and augmenting it with new capabilities.

With the advent of U-Net architectures, a wide range of medical imaging tasks has been implemented and produced state-of-the-art results since 2015 [20]. Recently, Chen *et al.* modified U-Net architecture by introducing sub-pixel layers to improve low-light imaging [3] and obtained promising results, with high signal-to-noise-ratio (SNR) and perfect color transformation on their own SID dataset. The authors in [10], [11] used nearest-neighbor interpolation for image reconstruction and super-resolution. The authors in [16] investigated the problem of transposed convolution and provided a solution by nearest-neighbor interpolation. However, the importance of integrating it into the deep CNN as part of the image upsampling operation was not fully explored so far. There have been a few papers tackling the segmentation and identification of surgical instruments from the endoscopic video image, and, even fewer than half a dozen papers tackling this challenge using deep learning. One notable research contribution has been the use of a modified version of FCN-8, yet with no attempts for multi-class segmentation [6].

Multi-class (both instrument part and type) tool segmentation was first proposed by Shvets *et al.* [24], and Pakhomov *et al.* [17] and achieved promising results. They modified the classic U-Net model [20] that relies on the transposed convolution or deconvolution, in a similar, yet opposite fashion to the convolutional layers. As an example, instead of mapping from 4×4 input pixels to 1 output pixel, they map from 1 input pixel to 4×4 output pixels. However, its computational performance is much slower, as the filters need additional weights and parameters that also require training in an end-to-end manner. Additionally, transposed convolution can easily lead to “uneven overlap”, characterized by checkerboard-like patterns resulting in artifacts on a variety of scales and colors [16]. Redford *et al.* [18] and Salimans *et al.* [21] introduced the drawback associated with those artifacts and checkerboard patterns generated by the transposed convolution. While it is difficult to entirely remove these limitations and their resulting artifacts, our goal is to, at first, minimize their occurrence.

Hence, in the efforts to mitigate these challenges associated with the classic U-Net architecture, in this paper, we present the U-NetPlus model by introducing both VGG-11 and

VGG-16 as an encoder with batch-normalized pre-trained weights and nearest-neighbor interpolation as the replacement of the transposed convolution in the decoder layer. This pre-trained encoder [9] speeds up convergence and leads to improved results by circumventing the optimization challenges associated with the target data [7]. Moreover, the nearest-neighbor interpolation used in the decoder section removes the artifacts generated by the transposed convolution.

To test the proposed U-NetPlus network, we implemented some of the recent state-of-the-art surgical tool segmentation architectures and compared their results to those of the U-NetPlus architecture. From the above-mentioned papers, only one seems to have achieved results comparable to ours [9], but it still suffers from several artifacts, which we have been able to further mitigate some of these artifacts using our proposed method. As such, while this paper leverages some of the existing infrastructures of fully convolutional network, it focuses on demonstrating the adaptation of existing infrastructure to refine its performance for a given task — in this case, the segmentation and identification of surgical instruments from endoscopic images — rather than proposing a new fully convolutional framework. We demonstrate that the potential use of nearest-neighbor interpolation in the decoder removes artifacts and reduces the number of parameters.

II. Methodology

A. Overview of Proposed Method

U-NetPlus has a downsampling path and an upsampling path, followed by a multi-class softmax layer for pixel-wise segmentation, as illustrated in Fig. 1. Similar to U-Net, our proposed U-NetPlus works like an auto-encoder with both a downsampling and an upsampling path. To maintain exactly the same number of channels as in the encoder part, downsampling and upsampling paths are connected through skip connections. This allows very precise alignment of the mask to the original image, which is particularly important in medical imaging. Furthermore, skip connections mitigate the vanishing gradient problem by initiating multiple paths for backpropagation. Generally, weights are initialized randomly to train a network. However, limited training data can introduce overfitting problems, which become very “expensive” as far as manually altering the segmentation mask. Therefore, transfer learning can be used to initialize the network weights. But since a surgical instrument is not a class of ImageNet, one way to use transfer learning for a new task is to partially reuse ImageNet feature extractor — VGG-11 or VGG-16 as encoder — and then add a decoder. An improvement has been introduced for the encoder part, where we initiated a pre-trained VGG-11 and VGG-16 architecture with batch-normalization layers that have 11 and 16 sequential layers, respectively. Following this modification, it has been shown the pre-trained model is able to train the network within a very short time and with greater accuracy [12].

The feature map of VGG-11 consists of seven convolutional layers of 3×3 kernel size followed by a ReLU activation function. For the reduction of the feature map size, max pooling with stride 1 was used. The number of channels is then doubled by the pooling operation, until reaching a total of 512 channels. Weights are copied from the original pre-trained VGG-11 on Imagenet.

The key effect of batch normalization has been investigated in a recent paper [22]. According to this work, batch normalization not only reduces the internal co-variate shift but also re-parameterizes the underlying gradient optimization problem that makes the training more predictive at a faster convergence. After analyzing the impact of inserting BatchNorm layer, we applied BatchNorm layer after each convolutional layer. The downsampling path decreases the feature size while increasing the number of feature maps, whereas the upsampling path increases the feature size while decreasing the number of feature maps, eventually leading to a pixel-wise mask. For the upsampling operation, we modified the existing architecture to reconstruct the high-resolution feature maps. Rather than using transposed convolution, we used the nearest-neighbor upsampling layer with a carefully selected stride and kernel size at the beginning of each block followed by two convolution layers and a ReLU function that would increase the spatial dimension in each block by a factor of 2.

Nearest-neighbor interpolation upsamples the input feature map by superimposing a regular grid onto it. Given I_i be the input grid which is to be sampled, the output grid is produced by a linear transformation $\tau_\theta(I_i)$. Therefore, for an upsampling operation, τ_θ can be defined as:

$$\begin{pmatrix} p_i^o \\ q_i^o \end{pmatrix} = \tau_\theta(I_i) = \begin{bmatrix} \theta & 0 \\ 0 & \theta \end{bmatrix} \begin{pmatrix} p_i^t \\ q_i^t \end{pmatrix}, \quad \theta \geq 1, \quad (1)$$

where $(p_i^o, q_i^o) \in I_i$ are the original sampling input coordinates, (p_i^t, q_i^t) are the target coordinates, and θ upsampling factor. The principle of how nearest-neighbor (NN) interpolation works to enlarge the image size, is shown in Fig. 1. After locating the center pixel of the cell of the output raster dataset on the input raster, the location of the nearest center of the cell on the input raster will be determined and the value of that cell on the output raster will be assigned afterward. As an example, we demonstrate the upsampling of a 4×4 image using this approach. The cell centers of the output raster are equidistant. A value is needed to be derived from the input raster for each output cell. Nearest-neighbor interpolation would select those cells centers from the input raster that are closest to that of the output raster. The black areas of the middle image can be filled with the copies of the center pixel. Therefore, this fixed interpolation weights requires no learning for upsampling operation compared to strided or transposed convolution leading to create a more memory efficient upsampling operation. The algorithm is similar to the one proposed and used by the authors of [4] in their work.

B. Dataset

For both training and validation, we used the Robotic instruments dataset from the sub-challenge of *MICCAI 2017 Endoscopic Vision Challenge* [1]. The training dataset has 8×225 frame sequences with 2 Hz frame rate of high-resolution stereo camera images collected from a da Vinci Xi surgical system during laparoscopic cholecystectomy procedures. The frames were re-sampled from 30 Hz video to 2 Hz to avoid any redundancy issues. A stereo camera was used to capture the video sequences that consists of the left and right eye views with a resolution of 1920×1080 in RGB format. In each frame, the surgical instrument was manually labeled by expert clinicians as a rigid shaft, wrist, and claspers.

The test set has 8×75 frame sequences and 2×300 frame videos. The challenge is to segment 7 classes such as grasping retractor, needle driver, prograsp forceps, vessel sealer, etc.

C. Data Augmentation

We augmented the MICCAI 2017 EndoVis Challenge data using the albumentations library that was reported as a fast and flexible implementation for data augmentation in [2]. These libraries include both affine and elastic transformations, and their effects on the image data during augmentation is illustrated in Fig. 2.

In short, the affine transformation includes scaling, translation, horizontal flip, vertical flip, random brightness and noise addition etc. For the elastic transformation (non-affine), first a random displacement field, $F(R)$ is generated for the horizontal and vertical directions, δx , and δy respectively where, $[\delta x, \delta y] = [-1 \quad \delta x, \delta y \quad +1]$.

These random fields are then convolved with an intermediate value of σ (in pixels) and the fields are multiplied by a scaling factor α that controls intensity. Thus, we obtain the elastically transformed image in which the global shape of the interest is undisturbed, unlike in the affine-transformed image.

D. Implementation Details

We implemented our methodology using PyTorch¹. During the pre-processing step, we cropped the un-wanted black border from each video frame. Images were normalized by subtracting their mean and dividing by their standard deviation (i.e., according to their z-scores). Batch normalization was used before each weighted layer, as it re-parameterizes the underlying gradient optimization problem that helps the training to converge faster [22]. For training, we used Adam optimizer as an optimizer with a learning rate of 0.00001. We didn't use dropout as it degraded validation performance in our case. All models were trained for 100 epochs. The training set was shuffled before each epoch and the batch size was 4 in our case. All experiments were run on a machine equipped with an NVIDIA GTX 1080 Ti GPU (11GBs of memory).

E. Performance Metrics

In this work, we used the common Jaccard index — also referred to as the intersection-over-union (IoU) — to evaluate segmentation results. It is an overlap index that quantifies the agreement between two segmented image regions: a ground truth segmentation and the predicted segmentation method. Given a vector of ground truth labels T_1 and a vector of predicted labels P_1 , IoU can be defined as

$$J(T_1, P_1) = \frac{|T_1 \cap P_1|}{|T_1 \cup P_1|} = \frac{|T_1 \cap P_1|}{|T_1| + |P_1| - |T_1 \cap P_1|}, \quad (2)$$

Eqn. 2 can further be clarified. Given a pixel j , the label of the pixel z_j and the probability of the same pixel for the predicted class $\hat{z}_j$, Eqn. 1 for k number of dataset

¹<https://github.com/pytorch/pytorch>

$$J = \frac{1}{k} \sum_{j=1}^k \left(\frac{z_j \hat{z}_j}{z_j + \hat{z}_j - z_j \hat{z}_j} \right), \quad (3)$$

We can represent the loss function in a common ground of *log* scale as this task is a pixel classification problem. So, for a given pixel j , the common loss can be defined as the function H for k number of dataset

$$H = -\frac{1}{k} \sum_{j=1}^k (z_j \log \hat{z}_j + (1 - z_j) \log(1 - \hat{z}_j)), \quad (4)$$

From both the Eqn. 1 and Eqn. 2, we can combine and can get a generalized loss

$$L = H - \log J \quad (5)$$

Our aim is to minimize the loss function, and, to do so, we can maximize the intersection, J between the predicted mask and the ground truth.

Another commonly used performance metric is the DICE coefficient. Given the set of all pixels in the image, set of foreground pixels by automated segmentation S_1^a , and the set of pixels for ground truth S_1^g , DICE score can be compared with $[S_1^a, S_1^g] \subseteq \Omega$, when a vector of ground truth labels T_1 and a vector of predicted labels P_1 ,

$$D(T_1, P_1) = \frac{2|T_1 \cap P_1|}{|T_1| + |P_1|} \quad (6)$$

DICE score will measure the similarity between two sets, T_1 and P_1 and $|T_1|$ denotes the cardinality of the set T_1 with the range of $D(T_1, P_1) \in [0, 1]$

III. Results

A. Quantitative Results

To illustrate the potential improvement in segmentation performance by using the nearest-neighbor interpolation (i.e., fixed upsampling) in the decoder, we conducted a paired comparison between the segmentation results obtained using the classical U-Net architecture, U-Net + NN, TernaNet, and U-NetPlus (our proposed method).

Training accuracy for binary segmentation is shown in Fig. 3 for 100 epochs. We compare our proposed architecture with three other models: U-Net, U-Net+NN, TernaNet. We can observe from this figure that after adding nearest-neighbor (NN) in the decoder of U-Net, the training accuracy of the classical U-Net framework (shown in blue) featuring the transposed convolution in the decodes, improves. Furthermore, the training of our proposed method (U-NetPlus) also converges faster and yields better training accuracy compared to TernaNet (shown in cyan). Hence, this graph alone illustrates the benefit of the nearest-neighbor interpolation on the segmentation performance.

The model was tested on the MICCAI 2017 EndoVis dataset. Table I summarizes the performance of our proposed U-NetPlus framework in the context of several state-of-the-art multi-task segmentation techniques. The table clearly indicates the improvement in segmentation following the addition of nearest-neighbor interpolation in the decoder step across all frameworks — U-Net and TerausNet. Moreover, our model had been compared with four different structures other than U-Net and TerausNet — ToolNetH, ToolNetMS, FCN-8s, and CSL. The last one (CSL) was the first approach to multi-class surgical instrument segmentation. But, they used only two instrument classes (shaft and clasps) and omit wrist class which we introduced in our approach and the overall accuracy that we obtained was significantly higher than the CSL approach.

We conducted a paired statistical test to compare the segmentation performance of each of these methods (U-Net, U-Net+NN, TerausNet, U-NetPlus) in terms of the IoU and DICE metric. As illustrated, our proposed U-NetPlus architecture yielded a statistically significant² 11.01% improvement ($p < 0.05$) in IoU and 6.91% DICE ($p < 0.05$) over the classical U-Net framework; a statistically significant 8.0% improvement ($p < 0.05$) in IoU and 5.79% DICE ($p < 0.05$) over the U-Net + NN framework; a statistically significant 0.18% improvement in IoU and 0.21% DICE ($p < 0.1$) over the state-of-the-art TerausNet framework [24].

Multi-class instrument segmentation was performed by labeling each instrument pixel with the corresponding index given in the training set. This application consisted of three classes: shaft, wrist, and clasps. The multi-class segmentation using our proposed U-NetPlus framework yielded a mean 65.75% IoU and 76.26% DICE. The accuracy and precision of the U-NetPlus architecture relative to the other three frameworks are illustrated in Fig. 3. As shown, the U-NetPlus framework outperforms the currently deemed best-in-class TerausNet framework.

The instrument type was segmented by labeling each instrument pixel with the corresponding instrument type, according to the training set, and all background pixels were labeled as 0. In the case of instrument type segmentation (for class = 7), U-NetPlus-VGG-11 encoder worked better than the U-NetPlus-VGG-16. Our results for instrument type segmentation can be further refined.

B. Qualitative Results

The qualitative comparison of our proposed model both for a binary and multi-class instrument segmentation is shown in Fig. 4. The second row of the figure shows that for the binary segmentation, the classical U-Net shows a portion of the instrument which was not present in the binary mask of our ground truth data (second row and second column). U-NetPlus shows the best performance for binary segmentation (i.e. it can clearly segment out the instruments from the background), whereas TerausNet still shows un-wanted regions in the segmentation output. For the instrument parts segmentation, U-Net still segments the un-wanted instrument (blue), whereas U-NetPlus can segment the 3 classes (blue:shaft, green:wrist, yellow:clasps) near perfectly compared to TerausNet. For the instrument

²For statistical significance testing, Wilcoxon signed-rank test is performed

type segmentation, we can clearly observe that U-Net can not differentiate between the blue and the green class, whereas U-NetPlus can differentiate these classes more accurately than TernaNet. Both the binary and multi-class segmented output have been overlaid onto the original image (sixth, seventh, eighth, and ninth column). The figure has a clear indication of qualitative improvement of U-NetPlus over U-Net, U-Net+NN, and TernaNet.

C. Attention Study

We performed an additional attention analysis to further analyze the segmentation performance. This attention study visualizes where our proposed algorithm “looks” in an image by using a novel image saliency technique [5] that learns the mask of an image by suppressing the softmax probability of its target class. Fig. 5 shows the heat-map image of the segmented surgical instruments superimposed onto the original video image.

Fig. 5 shows that the U-Net + NN architecture featuring the nearest-neighbor sampling in the decoder path and the traditional U-Net encoder out-performed the traditional U-Net architecture (featuring the transposed convolution in the decoder). On the other hand, due to the limited training dataset, the U-Net + NN framework slightly underperformed the TernaNet architecture featuring the pre-trained VGG network in the encoder. Nevertheless, using this class activation mapping, our proposed approach (U-NetPlus) localizes the wrist and clasps of the bipolar forceps near perfectly compared to the traditional U-Net, U-Net + NN, and TernaNet frameworks. Therefore, the skillful integration and combination of pre-trained encoder and nearest-neighbor interpolation as a fixed upsampling technique yields higher overall performance.

IV. Discussion and Conclusion

In this paper, we proposed a modified U-Net architecture for the surgical tool segmentation. To improve robustness beyond that of the U-Net framework, we used a pre-trained model as the encoder with batch-normalization, which converges much faster than the network trained from scratch. In the decoder part, we substituted the deconvolution layer with an upsampling layer that uses nearest-neighbor interpolation followed by two convolution layers. Moreover, we used a fast and effective data augmentation technique to avoid the overfitting problem. We evaluated its performance on the MICCAI 2017 EndoVis Challenge dataset. We also visualized the output of our proposed model both as stand-alone surgical instrument segmentation, as well as overlays onto the native endoscopic images. Apart from that, we also conducted an “attention” study to determine where our proposed algorithm “looks” in an image.

Our proposed model with batch-normalized U-NetPlus-VGG-16 outperforms existing methods according to both the Jaccard and DICE metrics, achieving 90.20% DICE for binary class segmentation and 76.26% for parts segmentation, both of which showed at least 0.21% improvement over the current methods and more than 6% improvement over the traditional U-Net architecture. Nevertheless, U-NetPlus-VGG-16’s performance with regards to identifying the instrument type was inferior to that of U-NetPlus-VGG-11, which was slightly superior to the other disseminated techniques. Though the improvement is still

small, our paired statistical test showed significant improvement over the performance of the state-of-the-art TerausNet method.

To evaluate the performance improvement in segmentation yielded by our proposed method, we conducted the above-mentioned paired statistical tests between the output of our proposed method and that of the other networks. The test showed that the U-NetPlus framework significantly out-performed the U-Net and TerausNet architectures ($p < 0.05$).

Although there are existing methods and approaches utilizing the interpolation on the upsampling path of an encoder-decoder network for different segmentation purposes, the masterly integration, and adaptation of existing methods for improving segmentation accuracy of the surgical instruments is a key of our research. Moreover, we emphasize our main contribution that lies in improving U-Net architecture via a modification of the state-of-the-art TerausNet to mitigate some of the artifacts still existing. So while this paper does not propose a fully novel framework, it does demonstrate that the skillful integration and combination of existing contributions yields higher overall performance.

Therefore, this paper serves as the first demonstration of a modified version of U-Net decoder via nearest-neighbor interpolation to remove artifacts induced by the transposed convolution. Our proposed architecture is used to segment the surgical instruments from laparoscopic images and showed improved performance over the state-of-the-art TerausNet framework.

Acknowledgments

This work was supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award No. R35GM128877 and by the Office of Advanced Cyber infrastructure of the National Science Foundation under Award No. 1808530.

References

- [1]. MICCAI 2017 Endoscopic Vision Challenge: Robotic Instrument Segmentation Sub-Challenge, 2017, <https://endovissub2017-roboticinstrumentsegmentation.grand-challenge.org/Data/>.
- [2]. Buslaev EKVIIA, Parinov A and Kalinin AA, "Albumentations: fast and flexible image augmentations," arXiv e-prints arXiv:1809.06839, 2018.
- [3]. Chen C, Chen Q, Xu J, and Koltun V, "Learning to see in the dark," arXiv preprint arXiv:1805.01934, 2018.
- [4]. Dong C, Loy CC, and Tang X, "Accelerating the super-resolution convolutional neural network," in European Conference on Computer Vision Springer, 2016, pp. 391–407.
- [5]. Fong RC and Vedaldi A, "Interpretable explanations of black boxes by meaningful perturbation," arXiv preprint arXiv:1704.03296, 2017.
- [6]. García-Peraza-Herrera LC, Li W, Fidon L, Gruijthuijsen C, Devreker A, Attilakos G, Deprest J, Vander Poorten E, Stoyanov D, Vercauteren T, et al., "Toolnet: holistically-nested real-time segmentation of robotic surgical tools," in Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on. IEEE, 2017, pp. 5717–5722.
- [7]. He K, Girshick R, and Dollár P, "Rethinking imagenet pre-training," arXiv preprint arXiv:1811.08883, 2018.
- [8]. He K, Gkioxari G, Dollár P, and Girshick R, "Mask r-cnn," in Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 2961–2969.
- [9]. Iglovikov V and Shvets A, "Terausnet: U-net with VGG11 encoder pre-trained on imagenet for image segmentation," arXiv e-prints arXiv:1801.05746, 2018.

- [10]. Jia X, Chang H, and Tuytelaars T, "Super-resolution with deep adaptive image resampling," arXiv preprint arXiv:1712.06463, 2017.
- [11]. Jiang N and Wang L, "Quantum image scaling using nearest neighbor interpolation," *Quantum Information Processing*, vol. 14, no. 5, pp. 1559–1571, 2015.
- [12]. Kornblith S, Shlens J, and Le QV, "Do better imagenet models transfer better?" arXiv preprint arXiv:1805.08974, 2018.
- [13]. Laina I, Rieke N, Rupperecht C, Vizcaíno JP, Eslami A, Tombari F, and Navab N, "Concurrent segmentation and localization for tracking of surgical instruments," in *International Conference on Medical Image Computing and Computer-Assisted Intervention Springer*, 2017, pp. 664–672.
- [14]. Long J, Shelhamer E, and Darrell T, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431–3440.
- [15]. Mechrez R, Goldberger J, and Greenspan H, "Patch-based segmentation with spatial consistency: application to MS lesions in brain MRI," *Journal of Biomedical Imaging*, vol. 2016, p. 3, 2016.
- [16]. Odena A, Dumoulin V, and Olah C, "Deconvolution and checkerboard artifacts," *Distill*, vol. 1, no. 10, p. e3, 2016.
- [17]. Pakhomov D, Premachandran V, Allan M, Azizian M, and Navab N, "Deep residual learning for instrument segmentation in robotic surgery," arXiv preprint arXiv:1703.08580, 2017.
- [18]. Radford A, Metz L, and Chintala S, "Unsupervised representation learning with deep convolutional generative adversarial networks," arXiv preprint arXiv:1511.06434, 2015.
- [19]. Rao PP, "Robotic surgery: new robots and finally some real competition!" *World Journal of Urology*, vol. 36, no. 4, pp. 537–541, 2018. [PubMed: 29427003]
- [20]. Ronneberger O, Fischer P, and Brox T, "U-Net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention Springer*, 2015, pp. 234–241.
- [21]. Salimans T, Goodfellow I, Zaremba W, Cheung V, Radford A, and Chen X, "Improved techniques for training GANs," in *Advances in Neural Information Processing Systems*, 2016, pp. 2234–2242.
- [22]. Santurkar S, Tsipras D, Ilyas A, and Madry A, "How does batch normalization help optimization? (no, it is not about internal covariate shift)," arXiv preprint arXiv:1805.11604, 2018.
- [23]. Shen D, Wu G, and Suk H-I, "Deep learning in medical image analysis," *Annual Review of Biomedical Engineering*, vol. 19, pp. 221–248, 2017.
- [24]. Shvets A, Rakhlin A, Kalinin AA, and Iglovikov V, "Automatic instrument segmentation in robot-assisted surgery using deep learning," arXiv preprint arXiv:1803.01207, 2018.
- [25]. Yang M, Yu K, Zhang C, Li Z, and Yang K, "DenseASPP for semantic segmentation in street scenes," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3684–3692.

Fig. 1.

(a) Modified U-Net with batch-normalized VGG11 as an encoder and upsampling as the decoder. Feature maps are denoted by rectangular shaped box. It consists of both an upsampling and a downsampling path and the feature map resolution is denoted by the box height, while the width represents the number of channels. Cyan arrows represent the max-pooling operation, whereas light-green arrows represent skip connections that transfer information from the encoder to the decoder. Red upward arrows represent the decoder which consists of nearest-neighbor upsampling with a scale factor of 2 followed by 2 convolution layers and a ReLU activation function; (b)-(d) working principle of nearest-neighbor interpolation where the low-resolution image is resized back to the original image.

Fig. 2.

Example images of applying both affine and elastic transformation in albumentations library for data augmentation.

Fig. 3. Quantitative comparison of (a) training accuracy (left), (b) multi-class (class=3) instrument parts (middle) (c) multi-task segmentation accuracy (right).

Fig. 4.

Qualitative comparison of binary segmentation, instrument part and instrument type segmentation result and their overlay onto the native endoscopic images of the MICCAI 2017 EndoVis video dataset yielded by four different frameworks: U-Net, U-Net+NN, TernausNet, and U-NetPlus.

Fig. 5.

Attention results: U-NetPlus “looks” at a focused target region, whereas U-Net, U-Net+NN and TerausNet appear less “focused”, leading to less accurate segmentation.

Quantitative comparison for instrument segmentation across several techniques. Mean and (standard deviation) values are reported for IoU(%) and DICE coefficient(%) from all networks against our proposed U-NetPlus. The statistical significance of the results for U-Net + NN and U-NetPlus model compared against the baseline model (U-Net and TernaNet) are represented by * and ** for p-values 0.1 and 0.05, respectively. U-Net has been compared with U-Net+NN, TernaNet has been compared with U-NetPlus. The best performance metric (IoU and DICE) in each category (Binary, Instrument Part and Instrument Type Segmentation) is indicated in **bold** text.

TABLE I

Network	Binary Segmentation			Instrument Part			Instrument Type		
Metric	IoU	DICE		IoU	DICE		IoU	DICE	
ToolNetH [6]	74.4	82.2		-	-		-	-	
ToolNetMS [6]	72.5	80.4		-	-		-	-	
FCN-8s [6]	70.9	78.8		-	-		-	-	
CSL [13]	-	88.9		-	87.70 (Shaft)		-	-	
U-Net [20]	75.44 (18.18)	84.37 (14.58)		48.41 (17.59)	60.75 (18.21)		15.80 (15.06)	23.59 (19.87)	
U-Net + NN	77.05** (15.71)	85.26* (13.08)		49.39* (15.18)	61.98* (15.47)		16.72* (13.45)	23.97 (18.08)	
TernaNet [24]	83.60 (15.83)	90.01 (12.50)		65.50 (17.22)	75.97 (16.21)		33.78 (19.16)	44.95 (22.89)	
U-NetPlus-VGG-11	81.32 (16.76)	88.27 (13.52)		62.51 (18.87)	74.57 (16.51)		34.84* (14.26)	46.07** (16.16)	
U-NetPlus-VGG-16	83.75 (13.36)	90.20* (11.77)		65.75 (14.74)	76.26* (13.54)		34.19 (15.06)	45.32 (17.86)	
					94.75(Shaft)				