

Towards Unsupervised Learning for Instrument Segmentation in Robotic Surgery with Cycle-Consistent Adversarial Networks

Daniil Pakhomov¹ Wei Shen¹ and Nassir Navab¹

Abstract—Surgical tool segmentation in endoscopic images is an important problem: it is a crucial step towards full instrument pose estimation and it is used for integration of pre- and intra-operative images into the endoscopic view. While many recent approaches based on convolutional neural networks have shown great results, a key barrier to progress lies in the acquisition of a large number of manually-annotated images which is necessary for an algorithm to generalize and work well in diverse surgical scenarios. Unlike the surgical image data itself, annotations are difficult to acquire and may be of variable quality. On the other hand, synthetic annotations can be automatically generated by using forward kinematic model of the robot and CAD models of tools by projecting them onto an image plane. Unfortunately, this model is very inaccurate and cannot be used for supervised learning of image segmentation models. Since generated annotations will not directly correspond to endoscopic images due to errors, we formulate the problem as an unpaired image-to-image translation where the goal is to learn the mapping between an input endoscopic image and a corresponding annotation using an adversarial model. Our approach allows to train image segmentation models without the need to acquire expensive annotations and can potentially exploit large unlabeled endoscopic image collection outside the annotated distributions of image/annotation data. We test our proposed method on Endovis 2017 challenge dataset and show that it is competitive with supervised segmentation methods.

I. INTRODUCTION

Robot-assisted Minimally Invasive Surgery (RMIS) provides a surgeon with improved control, facilitating procedures in confined and difficult to access anatomical regions. However, complications due to the reduced field-of-view provided by the surgical camera limit the surgeon's ability to self-localize. Computer assisted interventions (CAI) can help a surgeon by integrating additional information. For example, overlaying pre- and intra-operative imaging with the surgical console can provide a surgeon with valuable information which can improve decision making during complex procedures [25]. Integrating this data is a complex task and involves understanding relations between the patient anatomy, operating instruments and surgical camera. Segmentation of the instruments in the camera images is a crucial component of this process and can be used to prevent rendered overlays from occluding the instruments while providing crucial input to instrument tracking frameworks [18], [1].

Segmentation of surgical tools from tissue backgrounds is an extremely difficult task due to lighting challenges such as reflections, shadows and occlusions such as smoke and blood. Early methods attempted to simplify the problem by

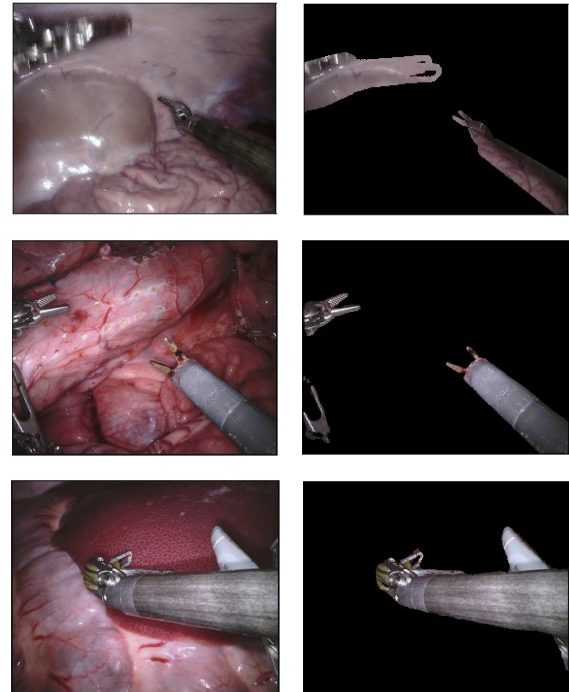


Fig. 1. Top row shows an example of synthetic annotation acquired using forward kinematics model that was used for training of our method. As it can be seen, it barely captures the actual tool due to errors. Two bottom rows show example segmentations delivered by our method on a random images from Endovis 2017 dataset that were not used for training.

modifying the appearance of the instruments [26]. However, this complicates clinical application of the technique as sterilization can become an issue. Segmentation of the instruments using natural appearance is a more desirable approach as it can be applied directly to pre-existing clinical setups. However, this defines a more challenging problem. To solve it, previous work has relied on machine learning techniques to model the complex discriminative boundary. Approaches based on Random Forests [3], maximum likelihood Gaussian Mixture Models [18] and Naive Bayesian classifiers [24], all trained on color features, have been applied. More recently, the state-of-the-art has increasingly been defined by Fully Convolutional Networks (FCNs), such as the FCN-8s model [15] adapted for the task of binary segmentation of robotic tools [7] and U-Net [22] which was used for both binary and instrument part segmentation [13]. In order for segmentation approaches based on supervised

¹Johns Hopkins University, Baltimore, USA

training of neural networks to successfully generalize, a considerable amount of annotated data is required [4]. At the same time, creating a segmentation dataset with high resolution fine annotations is an extremely time-consuming and costly process [4] [2].

Recently, unsupervised and self-supervised methods were introduced to make use of a great amount of unlabeled images of surgical instruments which can be collected with little or no effort: Ross et al. [23] proposed a method, which allows to pretrain segmentation models on recoloring task, greatly reducing the number of labeled images necessary for supervised learning; Rocha et al. [21] proposes an optimization method to obtain corrected labels to train a binary segmentation model despite imprecise kinematic model which sometimes results in poor results when the approximate labels are located far from the actual position of a surgical tool. Apart from instrument segmentation domain, Mahmood et al. [16] trained a model for multi-organ nuclei segmentation with synthetic and real data, Pfeiffer et al. [19] generated synthetic dataset for liver segmentation and showed promising results.

One way to acquire segmentation masks automatically is to use kinematic model of a robot and CAD model of the tools and project them onto the image plane of camera. Unfortunately, since the kinematic model is imprecise, the generated projections exhibit errors and cannot be directly used for supervised training. Although a lot of effort was made in improving supervised training of instrument segmentation models [17] [13] [7], a method that is able to successfully use this kind of data for training can potentially be supplied with unlimited amount of automatically generated data.

In this paper, we introduce an approach for binary surgical tool segmentation which does not need any manual annotation and only uses labels generated with imprecise kinematic model. The problem is posed as image-to-image translation [28] where we need to convert an image of surgical scene from RGB representation to semantic labels representation. Due to errors in the generated annotations that are caused by the imprecise kinematics, we are not able to train a segmentation in a supervised setting using acquired image/annotation pairs. Although we lack direct supervision in the form of image and annotation pairs, we employ set-level supervision: we are given a set of surgical images and a set of generated annotations which were created automatically without any manual annotation (see Fig. 4). By using this approach, we learn the mapping from surgical images to binary tool segmentation which is competitive with supervised algorithms that use expensive manually annotated data. This will allow segmentation models to be trained on larger amount of data and generalize better.

II. METHOD

A. Data Generation

In our work we are not using images with manually created labels similar to Endovis 2017 [2] for training and, instead, rely only on labels generated with imprecise kinematics. We record image sequences of surgical procedures and

corresponding annotation masks. The masks are generated by rendering CAD models of each instrument that is attached to a da Vinci Xi system using joint encoder values that are synchronized to the video feed. Rigid body transforms between the robot base frame and the instruments and the camera are computed using forward kinematics from the DH parameters of the robot. Camera calibration is also acquired from the system. Inaccuracies exist between the true instrument and camera poses due to unknown hand-eye calibration transforms and errors from slack and tension in the cable driven arms of the da Vinci [1].

Overall, three video sequences of different procedures were collected resulting in 6 thousand frames and annotations. As mentioned previously, annotations exhibit errors similar to example presented in Fig. 1. The process of the data creation is fully automatic and does not involve any manual annotation and, therefore, can be used to generate a dataset of potentially unlimited size (see Fig. 4).

Collected video sequences contain 3 different types of robotic surgical instruments: Large Needle Driver, Prograsp Forceps, Bipolar Forceps. An example image sample from the dataset is shown in Fig. 1. Although we were not able to collect video sequences featuring all instruments [2] because their CAD models are not available, we noticed that our approach successfully generalizes to previously unseen instruments.

B. Set-Level Supervision

Since the annotations that were automatically generated for our collected images contain errors, we can not use supervision on the level of image/annotation pairs during training. Instead, we use set-level supervision: we propose an unpaired GAN-based approach that learns mappings between surgical images (I) and annotation masks (A) and enforces them to be cycle consistent. We use cycleGAN [28] approach to learn a mapping $G_A : I \rightarrow A$ between surgical images and corresponding annotation masks that generalizes to previously unseen surgical images.

Overall, our approach employs four networks: G_I (segmentation labels to surgical image generator), G_A (surgical image to segmentation labels generator), D_A (discriminator network of G_A), and D_I (discriminator network of G_I). The final objective [28] consists of two adversarial loss terms \mathcal{L}_{GAN} and cycle consistency loss term \mathcal{L}_{cyc} . After the training is done, we use G_A as our segmentation model and evaluate it on previously unseen images to assess its segmentation performance.

First, we introduce an adversarial loss [28] that forces our segmentation model G_A to translate surgical images to realistic segmentation masks by making them look similar to annotations that we collected from kinematics model:

$$\mathcal{L}_{GAN}(G_A, D_A) = \mathbb{E}_{a \sim p_{data}(a)} [\log D_A(a)] + \mathbb{E}_{i \sim p_{data}(i)} [\log(1 - D_A(G_A(i)))] \quad (1)$$

Where D_A is a discriminator network that learns to distinguish between segmentations generated by our model G_A

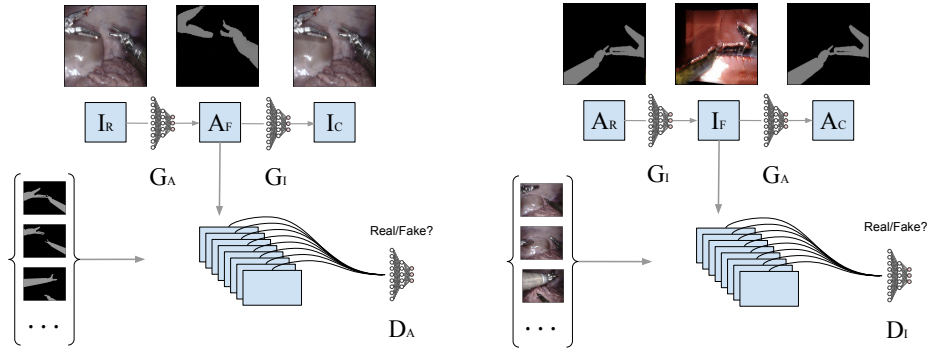


Fig. 2. Figure demonstrates main stages of our approach. During training, we learn four mappings which are represented by convolutional neural networks: generators G_A and G_I which map images to annotations and annotations to images respectively; discriminators D_A and D_I which learn to differentiate between real and generated annotations and images respectively. I_R and A_R stand for real images and annotations. I_F and A_F stand for images and annotations generated by networks. I_C and A_C stand for images and annotations generated from A_F and I_F respectively. They are used in cycle-consistency loss term. Loss also includes adversarial terms that forces discriminator to differentiate between real samples and the ones generated by the generators, and the generators to fool the discriminators. After the training, we use G_A to acquire segmentations of unseen images.

and real segmentations. Distributions $p_{\text{data}}(i)$ and $p_{\text{data}}(a)$ are approximated by surgical images and annotation samples collected in the previous step using kinematics model. However, optimizing this objective alone does not guarantee that a meaningful mapping G_A will be learnt since there are infinitely many mappings that will satisfy this criteria [28].

To make our problem more well-defined we follow [28] and add a constraint on translation to be cycle-consistent: we introduce a reverse mapping $G_I: A \rightarrow I$ from surgical images to corresponding segmentation masks (Fig. 2) and make G_I and G_A to be inverses of each other. Note, that the reverse mapping is introduced only to train our segmentation model G_A more effectively [28].

Second adversarial term of loss is introduced to force the distribution of surgical images generated by the network G_I to be similar to that of the real surgical images:

$$\mathcal{L}_{\text{GAN}}(G_I, D_I) = \mathbb{E}_{i \sim p_{\text{data}}(i)} [\log D_I(i)] + \mathbb{E}_{a \sim p_{\text{data}}(a)} [\log(1 - D_I(G_I(a)))] \quad (2)$$

where the mapping G_I is constrained to generate surgical images from real segmentation masks so that they look like real surgical images. Discriminator D_I is trying to differentiate between real and generated surgical images. This term introduces competition between generator and discriminator networks, allowing them to improve during training as a result of the competition [28].

As previously mentioned, in order to make our mapping more well-defined and learn a better segmentation model we also introduce a cycle consistency loss:

$$\mathcal{L}_{\text{cyc}}(G_A, G_I) = \mathbb{E}_{i \sim p_{\text{data}}(i)} [\|G_I(G_A(i)) - i\|_1] + \mathbb{E}_{a \sim p_{\text{data}}(a)} [\|G_A(G_I(a)) - a\|_1] \quad (3)$$

This term adds an additional regularization by making G_I and G_A inverse functions of each other. Specifically, it

ensures that $G_I(G_A(i)) \approx i$ (forward cycle consistency) and $G_A(G_I(a)) \approx a$ (backwards cycle consistency).

The full objective for our set-level supervised learning can be written as:

$$\mathcal{L}(G_A, G_I, D_A, D_I) = \mathcal{L}_{\text{GAN}}(G_A, D_A) + \mathcal{L}_{\text{GAN}}(G_I, D_I) + \lambda \mathcal{L}_{\text{cyc}}(G_A, G_I) \quad (4)$$

Where λ is a hyperparameter. We aim to solve:

$$G_A^*, G_I^* = \arg \min_{G_A, G_I} \max_{D_A, D_I} \mathcal{L}(G_A, G_I, D_A, D_I) \quad (5)$$

C. Edge Consistency

Applying the aforementioned method as-is allows us to learn the mapping G_A that acts as a segmentation model and delivers realistically looking segmentation masks but upon inspection, they are completely unaligned with instruments present in the input image. Intersection over union accuracy measure is also very low, which means that the segmentation has little or no overlap with tools located in the image. This motivates us to introduce additional constraints so that generated annotations are more aligned with instruments located in the surgical images.

We replace U-Net [22] architecture that we used for generator networks G_I and G_A with a network based on deep residual connections which give us much better results [9] [28]. We hypothesize that an implicit regularization that residual connections provide [8] allows us to acquire segmentation results that are more aligned with the input image.

However, the problem is not completely solved: while the intersection over union score is good, the predictions delivered by the method are not precise along the borders (See Fig. 3). In order for the segmentation method to be useful for augmented reality applications and tool tracking, the segmentation results should be aligned with the borders



Fig. 3. An example of segmentation results delivered by a deep residual model without edge consistency loss term. As it can be seen, while the shape of the segmentation looks realistic, it is not aligned with the edges of the tools. Since the segmentation result still has a great overlap with the tools, the intersection over union score is good but applications involving augmented reality and tracking of tools require the segmentation algorithm to be more precise along the borders.




Fig. 4. Supervised learning approach typically uses paired training data (left) consisting of training examples where every image has a corresponding manually annotated image. These datasets are expensive to collect and are usually of a small size. We instead consider unpaired training data (right), consisting of images and synthetically generated annotations using forward kinematics. This data can be automatically generated during medical procedures and can potentially be of unlimited size $N \gg n$.

of the actual tool [2]. Ideally, we would want to have consistent edges in the image and generated annotation (See Fig. 5). Inspired by a similar problem in the field of image matting [14] [20], we add the edge consistency term [14] to our loss:

$$\mathcal{L}_{\text{edge}}(G_A) = \mathbb{E}_{i \sim p_{\text{data}}(i)} [\mathcal{L}_C(G_A(i), i)], \quad (6)$$

$$\mathcal{L}_C(A, I) = \frac{\sum A_{\text{mag}} [1 - (I_x A_x + I_y A_y)^2]}{\sum A_{\text{mag}}} \quad (7)$$

where (I_x, I_y) and (A_x, A_y) are the normalized image and annotation gradients, and A_{mag} is the annotation gradient magnitude. Intuitively, the loss constrains the generator G_A to deliver segmentation masks that are more aligned with the image edges (See Fig. 5). Using both the deep residual network and edge consistency gives the best results (See Fig. 1).

III. EXPERIMENTS AND RESULTS

A. Implementation Details

1) *Network Architectures*: The generator architecture G_I and our segmentation model G_A are both based on deep residual network architecture [9], since it was shown to work better for our task. First two layers of the networks




Fig. 5. Examples of surgical image and annotation and their gradient magnitude images. The figures serves as a motivation for edge consistency loss: edges of the image should be aligned with the edges of the segmentation generated by the network. (Image is best viewed in the electronic version of the document)

subsample the input image by a factor of 4, followed by nine residual blocks and two upsampling layers with learnt filters that bring the output to the same size as the input image. The discriminator networks D_I and D_A have much simpler architectures based on PatchGANs [11] with three layers and fewer parameters, as suggested in [28]. All networks are implemented in a fully-convolutional fashion, which allows them to be applied to images of varying sizes [15].

In order to reduce memory consumption and be able to store four networks in GPU memory at the same time during training we had to resort to training with batch size one. Since training with small batch size with batch normalization is known to be unstable [10], we are employing instance normalization layers instead [27].

2) *Training details*: All the networks in our work were trained from scratch, starting with randomly initialized weights. We perform the optimization with Adam optimizer [12] with batch size one and learning rate of 0.0002. Overall, we train for 20 epochs with fixed learning rate and then linearly decay it to zero for another 20 epochs.

B. Accuracy measures

There are three commonly used accuracy measures for assessment of image segmentation models [5]:

- 1) Overall Pixel accuracy
- 2) Per-Class accuracy
- 3) Jaccard Index (Intersection-over-union)

In order to describe accuracy measures we define confusion matrix \mathbf{C} , which contains predictions for the whole image segmentation dataset \mathbf{D} :

$$C_{ij} = \sum_{z \in \mathbf{D}} |z \in \mathbf{I} \text{ such that } S_{gr}^I(z) = i \text{ and } S_{ps}^I(z) = j|$$

Where $S_{gr}^I(z)$ is the ground-truth label of pixel z in image \mathbf{I} , $S_{ps}^I(z)$ is a label predicted by a particular algorithm and $|\mathbf{A}|$ is the cardinality of the set \mathbf{A} . We define $\mathbf{G}_i = \sum_{j=1}^L C_{ij}$, where L is the number of classes and $\mathbf{P}_j = \sum_i C_{ij}$

Overall Pixel accuracy measures number of correctly classified pixels:

$$OP = \frac{\sum_{i=1}^L C_{ii}}{\sum_{i=1}^L G_i}$$

One significant limitation of this measure is its bias in the presence of very imbalanced classes [5]. If a dataset has one class that is more present than others and a segmentation model classifies it correctly, while making mistakes on other smaller classes, the value of the measure will not sufficiently represent that.

Per-Class accuracy does the same measurement as Overall Pixel accuracy but solves its problem with unbalanced classes by scaling results of from each class [5]:

$$PC = \frac{1}{L} \sum_{i=1}^L \frac{C_{ii}}{G_i}$$

While it solves the problem, another weakness can be observed: if a large background class is present, one achieves a better score by labeling object classes correctly while making more errors in labeling the background class [5].

Jaccard Index (Intersection-over-union) measures intersection over union for each class and reports the average between all classes:

$$JI = \frac{1}{L} \sum_{i=1}^L \frac{C_{ii}}{G_i + P_i - C_{ii}}$$

It solves the problem of previous two measures and currently is a main measure of semantic segmentation accuracy for PASCAL VOC challenge [6] and Endovis 2017 Robotic Instrument Segmentation Challenge [2] which we use to assess our method. In order to be able to compare our method with other methods on Endovis 2017 challenge we report Intersection over Union accuracy measure of our method on test dataset.

C. Dataset

We test our method on the EndoVis 2017 Robotic Instruments dataset [2]. There are 8 high resolution (1280×1024) sequences with 225 frames each in the training dataset [2]. As mentioned previously, we did not use the training dataset of Endovis 2017 and, instead, trained our method on our data acquired with imprecise kinematics model. Each pixel is labeled as either tool or background. There are 10 75-frame sequences in the test dataset that features 7 different robotic surgical instruments [2]. Samples from the dataset and segmentations delivered by our algorithm are depicted in Fig. 1.

The dataset is very challenging and even some of the supervised image segmentation methods were struggling to achieve good performance as it can be seen from the Table. I. At the same time it provides a good indication of whether our method can generalize well to unseen surgical scenes and instruments. After training our segmentation model on collected images and annotations generated with imprecise kinematics model, we evaluated it on the test set of the Endovis dataset.

D. Quantitative Study and Results

We used data from ten different test sequences of Endovis 2017 dataset to evaluate our trained instrument segmentation network. To assess generalizability of the developed algorithm, we also payed attention on how our method performed segmentation with instruments that were not represented in our collected dataset: our method successfully segmented previously unseen instrument (See bottom row of Fig. 1). Table I summarizes the quantitative results of the testing with intersection-over-union metric.

A comparative analysis with supervised segmentation methods that participated in the challenge was performed and our method, as can be seen from the table, outperforms four out of ten methods. We are the first ones to test a self-supervised method on challenging Endovis 2017 competition while outperforming some of the supervised methods that reported their results.

IV. CONCLUSION AND FUTURE WORK

Automated training of accurate instrument segmentation models for surgical procedures has the potential to completely eliminate costs associated with manual creation of datasets and can greatly affect the field by improving the segmentation performance and robustness of segmentation models by employing abundant unlabeled data. In this work we address the problem of training a segmentation model without direct supervision where images and inaccurate labels are generated automatically, therefore, eliminating the need for dataset creation.

We propose an approach that allows the instrument segmentation network to be trained on images with synthetically generated annotations with errors. The problem is posed as an unpaired image-to-image translation task. This way we are able to enforce set-level supervision between sets of surgical images and annotations. This approach performs on par with some standard supervised approaches tested on challenging Endovis 2017 dataset.

In the future work, we plan on adapting this approach to multi-class instrument segmentation and instance segmentation since this data can be easily automatically generated in a similar way to how we generated binary masks. Potentially, other types of mappings can also be learnt without direct supervision that can be very useful for surgical scene analysis and pose estimation like pixel-wise depth estimation and surgical tools landmarks detection. On the other hand, since our approach still performs worse than some of the supervised methods in terms of accuracy, a better segmentation network architecture can be used to close this gap in accuracy and completely eliminate the need for manual segmentation dataset creation.

REFERENCES

- [1] Max Allan, Ping-Lin Chang, Sébastien Ourselin, David J Hawkes, Ashwin Sridhar, John Kelly, and Danail Stoyanov. Image based surgical instrument pose estimation with multi-class labelling and optical flow. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 331–338. Springer, 2015.

	NCT	UB	BIT	MIT	SIAT	UCL	TUM	Delhi	UA	UW	Ours	+Edge
Dataset 1	0.784	0.807	0.275	0.854	0.625	0.631	0.760	0.408	0.413	0.337	0.692	0.727
Dataset 2	0.788	0.806	0.282	0.794	0.669	0.645	0.799	0.524	0.463	0.289	0.735	0.769
Dataset 3	0.926	0.914	0.455	0.949	0.897	0.895	0.916	0.743	0.703	0.483	0.721	0.755
Dataset 4	0.934	0.925	0.310	0.949	0.907	0.883	0.915	0.782	0.751	0.678	0.752	0.782
Dataset 5	0.701	0.740	0.220	0.862	0.604	0.719	0.810	0.528	0.375	0.219	0.778	0.794
Dataset 6	0.876	0.890	0.338	0.922	0.843	0.852	0.873	0.292	0.667	0.619	0.793	0.815
Dataset 7	0.846	0.930	0.404	0.856	0.832	0.710	0.844	0.593	0.362	0.325	0.686	0.694
Dataset 8	0.881	0.904	0.366	0.937	0.513	0.517	0.895	0.562	0.797	0.506	0.787	0.815
Dataset 9	0.789	0.855	0.236	0.865	0.839	0.808	0.877	0.626	0.539	0.377	0.673	0.726
Dataset 10	0.899	0.917	0.403	0.905	0.899	0.869	0.909	0.715	0.689	0.603	0.706	0.727
Mean IOU	0.843	0.875	0.326	0.888	0.803	0.785	0.873	0.612	0.591	0.461	0.732	0.760

TABLE I

QUANTITATIVE RESULTS OF OUR METHOD AND COMPARISON WITH SUPERVISED METHODS IN BINARY SEGMENTATION OF ROBOTIC TOOLS [2].

- [2] Max Allan, Alex Shvets, Thomas Kurmann, Zichen Zhang, Rahul Duggal, Yun-Hsuan Su, Nicola Rieke, Iro Laina, Niveditha Kalavakonda, Sebastian Bodenstedt, et al. 2017 robotic instrument segmentation challenge. *arXiv preprint arXiv:1902.06426*, 2019.
- [3] David Bouget, Rodrigo Benenson, Mohamed Omran, Laurent Riffaud, Bernt Schiele, and Pierre Jannin. Detecting surgical tools by modelling local appearance and global shape. *IEEE transactions on medical imaging*, 34(12):2603–2617, 2015.
- [4] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [5] Gabriela Csurka, Diane Larlus, Florent Perronnin, and France Meylan. What is a good evaluation measure for semantic segmentation?. In *BMVC*, volume 27, page 2013. Citeseer, 2013.
- [6] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [7] L. C. Garcia-Peraza-Herrera, W. Li, C. Gruijthuijsen, A. Devreker, G. Attilakos, J. Deprest, E. Vander Poorten, D. Stoyanov, T. Vercauteren, and S. Ourselin. Real-time segmentation of non-rigid surgical tools based on deep learning and tracking. In *CARE Workshop (MICCAI)*, 2016.
- [8] Klaus Greff, Rupesh K Srivastava, and Jürgen Schmidhuber. Highway and residual networks learn unrolled iterative estimation. *arXiv preprint arXiv:1612.07771*, 2016.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [10] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [11] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [12] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] Iro Laina, Nicola Rieke, Christian Rupprecht, Josué Page Vizcaíno, Abouzar Eslami, Federico Tombari, and Nassir Navab. Concurrent segmentation and localization for tracking of surgical instruments. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 664–672, 2017.
- [14] Alex Levinshtein, Cheng Chang, Edmund Phung, Irina Kezele, Wenzhangzhi Guo, and Parham Aarabi. Real-time deep hair matting on mobile devices. In *2018 15th Conference on Computer and Robot Vision (CRV)*, pages 1–7. IEEE, 2018.
- [15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015.
- [16] Faisal Mahmood, Daniel Borders, Richard Chen, Gregory N McKay, Kevan J Salimian, Alexander Baras, and Nicholas J Durr. Deep adversarial training for multi-organ nuclei segmentation in histopathology images. *IEEE transactions on medical imaging*, 2019.
- [17] Daniil Pakhomov, Vittal Premachandran, Max Allan, Mahdi Azizian, and Nassir Navab. Deep residual learning for instrument segmentation in robotic surgery. In *International Workshop on Machine Learning in Medical Imaging*, pages 566–573. Springer, 2019.
- [18] Zachary Pezzementi, Sandrine Voros, and Gregory D Hager. Articulated object tracking by rendering consistent appearance parts. In *Robotics and Automation, 2009. ICR'A'09. IEEE International Conference on*, pages 3940–3947. IEEE, 2009.
- [19] Micha Pfeiffer, Isabel Funke, Maria R Robu, Sebastian Bodenstedt, Leon Strenger, Sandy Engelhardt, Tobias Roß, Matthew J Clarkson, Kurinchi Gurusamy, Brian R Davidson, et al. Generating large labeled data sets for laparoscopic image processing tasks using unpaired image-to-image translation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 119–127. Springer, 2019.
- [20] Christoph Rhemann, Carsten Rother, Jue Wang, Margrit Gelautz, Pushmeet Kohli, and Pamela Rott. A perceptually motivated online benchmark for image matting. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1826–1833. IEEE, 2009.
- [21] Cristian da Costa Rocha, Nicolas Padoy, and Benoit Rosa. Self-supervised surgical tool segmentation using kinematic information. *arXiv preprint arXiv:1902.04810*, 2019.
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015.
- [23] Tobias Ross, David Zimmerer, Anant Vemuri, Fabian Isensee, Manuel Wiesenfath, Sebastian Bodenstedt, Fabian Both, Philip Kessler, Martin Wagner, Beat Müller, et al. Exploiting the potential of unlabeled endoscopic video data with self-supervised learning. *International journal of computer assisted radiology and surgery*, 13(6):925–933, 2018.
- [24] Stefanie Speidel, Michael Delles, Carsten Gutt, and Rüdiger Dillmann. Tracking of instruments in minimally invasive surgery for surgical skill analysis. In *International Workshop on Medical Imaging and Virtual Reality*, pages 148–155. Springer, 2006.
- [25] Russell H Taylor, Arianna Menciassi, Gabor Fichtinger, and Paolo Dario. Medical robotics and computer-integrated surgery. In *Springer handbook of robotics*, pages 1199–1222. Springer, 2008.
- [26] Oliver Tonet, TU Ramesh, Giuseppe Megali, and Paolo Dario. Tracking endoscopic instruments without localizer: image analysis-based approach. *Studies in health technology and informatics*, 119:544–549, 2005.
- [27] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.
- [28] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.