



Rethinking Surgical Instrument Segmentation: A Background Image Can Be All You Need

An Wang^{1,2}, Mobarakol Islam³, Mengya Xu⁴, and Hongliang Ren^{1,2,4}(✉)

¹ Department of Electronic Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong SAR, China

wa09@link.cuhk.edu.hk, hlren@ee.cuhk.edu.hk

² Shun Hing Institute of Advanced Engineering, The CUHK, Shatin, Hong Kong SAR, China

³ BioMedIA Group, Department of Computing, Imperial College London, London, UK

m.islam20@imperial.ac.uk

⁴ Department of Biomedical Engineering, National University of Singapore, Singapore, Singapore

mengya@u.nus.edu

Abstract. Data diversity and volume are crucial to the success of training deep learning models, while in the medical imaging field, the difficulty and cost of data collection and annotation are especially huge. Specifically in robotic surgery, data scarcity and imbalance have heavily affected the model accuracy and limited the design and deployment of deep learning-based surgical applications such as surgical instrument segmentation. Considering this, we rethink the surgical instrument segmentation task and propose a one-to-many data generation solution that gets rid of the complicated and expensive process of data collection and annotation from robotic surgery. In our method, we only utilize a single surgical background tissue image and a few open-source instrument images as the seed images and apply multiple augmentations and blending techniques to synthesize amounts of image variations. In addition, we also introduce the chained augmentation mixing during training to further enhance the data diversities. The proposed approach is evaluated on the real datasets of the EndoVis-2018 and EndoVis-2017 surgical scene segmentation. Our empirical analysis suggests that without the high cost of data collection and annotation, we can achieve decent surgical instrument segmentation performance. Moreover, we also observe that our method can deal with novel instrument prediction in the deployment domain. We hope our inspiring results will encourage researchers to emphasize data-centric methods to overcome demanding deep learning limitations besides data shortage, such as class imbalance, domain adaptation, and incremental learning. Our code is available at https://github.com/lofrienger/Single_SurgicalScene_For_Segmentation.

A. Wang and M. Islam—Co-first authors.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022

L. Wang (Eds.): MICCAI 2022, LNCS 13437, pp. 355–364, 2022.

https://doi.org/10.1007/978-3-031-16449-1_34

1 Introduction

Ever-larger models processing larger volumes of data have propelled the extraordinary performance of deep learning-based image segmentation models in recent decades, but obtaining well-annotated and perfectly-sized data, particularly in the medical imaging field, has always been a great challenge [6]. Various causes, including tremendous human efforts, unavailability of rare disease data, patient privacy concerns, high prices, and data shifts between different medical sites, have made acquiring abundant high-quality medical data a costly endeavor. Besides, dataset imperfection like class imbalance, sparse annotations, noisy annotations and incremental-class in deployment [20] also affects the training and deployment of deep learning models. Moreover, for the recent-developed surgery procedures like the single-port robotic surgery where no dataset of the new instruments is available [5], the segmentation task can hardly be accomplished. In the presence of these barriers, one effective solution to overcome the data scarcity problems is to train with a synthetic dataset instead of a real one.

A few recent studies utilize synthetic data for training and achieve similar and even superior performance than training with real data. For example, in the computer vision community, Tremblay et al. [19] develop an object detection system relying on domain randomization where pose, lighting, and object textures are randomized in a non-realistic manner; Gabriel et al. [7] make use of multiple generative adversarial networks (GANs) to improve data diversity and avoid severe over-fitting compared with a single GAN; Kishore et al. [14] propose imitation training as a synthetic data generation guideline to introduce more underrepresented items and equalize the data distribution to handle corner instances and tackle long-tail problems.

In medical applications, many works have focused on GAN-based data synthesizing [3, 10, 11, 18], while a few works utilize image blending or image composition to generate new samples. For example, mix-blend [8] mixes several synthetic images generated with multiple blending techniques to create new training samples. Nonetheless, one limitation of their work is that they need to manually capture and collect thousands of foreground instrument images and background tissue images, making the data generation process trivial and time-consuming. In addition, E. Colleoni et al. [4] recorded kinematic data as the data source to synthesize a new dataset for the instrument - Large Needle Drivers. In comparison with previous works, our approach only utilizes a single background image and dozens of foreground instrument images as the data source. Without costly data collection and annotation, we show the simplicity and efficacy of our dataset generation framework.

Contributions. In this work, we rethink the surgical instrument segmentation task from a data-centric perspective. Our contributions can be summarized as follows:

- With minimal human effort in data collection and without manual image annotations, we propose a data-efficient framework to generate high-quality synthetic datasets used for surgical instrument segmentation.
- By introducing various augmentation and blending combinations to the foreground and background source images, and training-time chained augmentation mixing, we manage to increase the data diversity and balance the instruments class distribution.
- We evaluate our method on two real datasets. The results suggest that our dataset generation framework is simple yet efficient. It is possible to achieve acceptable surgical instrument segmentation performance, even for novel instruments, by training with synthetic data that only employs a single surgical background image.

2 Proposed Method

2.1 Preliminaries

Data augmentation has become a popular strategy for boosting the size of a training dataset to overcome the data-hungry problem when training the deep learning models. Besides, data augmentation can also be regarded as a regularisation approach for lowering the model generalization error [9]. In other words, it helps boost performance when the model is tested on a distinct unseen dataset during training. Moreover, the class imbalance issue, commonly seen in most surgical datasets, can also be alleviated by generating additional data for the under-represented classes.

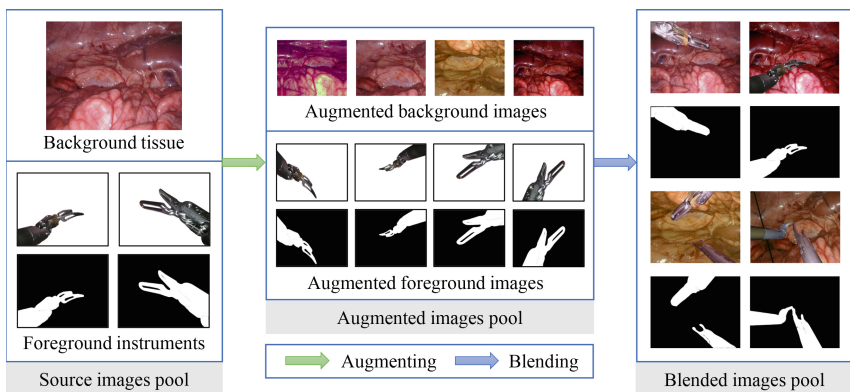


Fig. 1. Demonstration of the proposed dataset generation framework with augmenting and blending. With minimal effort in preparing the source images, our method can produce large amounts of high-quality training samples for the surgical segmentation task.

Blending is a simple yet effective way to create new images simply by image mixing or image composition. It can also be treated as another kind of data augmentation technique that mixes the information contained in different images instead of introducing invariance to one single image. Denote the foreground image and background image as x_f and x_b , we can express the blended image with a blending function Θ as

$$x = \Theta(x_f, x_b) = x_f \oplus x_b \quad (1)$$

where \oplus stands for pixel-wise fusion.

Training-time augmentation can help diversify training samples. By mixing various chained augmentations with the original image, more image variations can be created without deviating too far from the original image, as proposed by AugMix [12]. In addition, intentionally controlling the choices of augmentation operations can also avoid hurting the model due to extremely heavy augmentations. A list of augmentation operations is included in the augmentation chains, such as auto-contrast, equalization, posterization, solarization, etc.

2.2 Synthesizing Surgical Scenes from a Single Background

Background Tissue Image Processing. We collect one background tissue image from the open-source EndoVis-2018 dataset¹ where the surgical scene is the nephrectomy procedures. The critical criterion of this surgical background selection is that the appearance of the instrument should be kept as little as possible. In the binary instrument segmentation task, the background pixels are all assigned with the value 0. Therefore, the appearance of instruments in the source background image will occupy additional effort to handle. Various augmentations have been applied to this single background source image with the `imgaug`² library [13], including `LinearContrast`, `FrequencyNoiseAlpha`, `AddToHueAndSaturation`, `Multiply`, `PerspectiveTransform`, `Cutout`, `Affine`, `Flip`, `Sharpen`, `Emboss`, `SimplexNoiseAlpha`, `AdditiveGaussianNoise`, `CoarseDropout`, `GaussianBlur`, `MedianBlur`, etc. We denote the generated p variations of the background image as the background images pool $X_b^p = \{x_b^1, x_b^2, \dots, x_b^p\}$. As shown in Fig. 1, various augmented background images are generated from the single source background tissue image to cover a wide range of background distribution.

Foreground Instruments Images Processing. We utilize the publicly available EndoVis-2018 [1] dataset as the open resource to collect the seed foreground images. There are 8 types of instruments in the EndoVis-2018 [1] dataset, namely Maryland Bipolar Forceps, Fenestrated Bipolar instruments, Prograsp Forceps, Large Needle Driver, Monopolar Curved Scissors, Ultrasound Probe, Clip Applier, and Suction Instrument. We only employ 2 or 3 images for each

¹ <https://endovissub2018-roboticscenesegmentation.grand-challenge.org/>.

² <https://github.com/aleju/imgaug>.

instrument as the source images. We extract the instruments and make their background transparent. The source images are selected with prior human knowledge of the target scenes to ensure their high quality. For example, for some instruments like Monopolar Curved Scissors, the tip states (open or close) are crucial in recognition, and they are not reproducible simply by data augmentation. Therefore, we intentionally select source images for such instruments to make it possible to cover different postures and states. In this way, we aim to increase the in-distribution data diversity to substantially improve generalization to out-of-distribution (OOD) category-viewpoint combinations [15]. Since we get rid of annotation, the instrument masks are applied with the same augmentations as the instruments to maintain the segmentation accuracy. We denote the generated q variations of the foreground images as the foreground image pool $X_f^q = \{x_f^1, x_f^2, \dots, x_f^q\}$. Figure 1 shows some new synthetic instruments images. The foreground images pool, together with the background images pool, forms the augmented images pool, which is used for the following blending process.

Blending Images. After obtaining the background image pool X_b^p and the foreground image pool X_f^q , we randomly draw one sample from these two pools and blend them to form a new composited image. Specifically, the foreground image is pasted on the background image with pixel values at the overlapped position taken from the instruments. Furthermore, considering the real surgical scenes, the number of instruments in each image is not fixed. We also paste two instrument images on the background occasionally. Due to this design, we expect the model could better estimate the pixel occupation of the instruments in the whole image. Denoting the blended image as x_s , finally, the blended images pool with t synthetic images can be presented as $X_s^t = \{x_s^1, x_s^2, \dots, x_s^t\} = \{\Theta(x_f^i, x_b^j)\}$, where $i = 1, 2, \dots, p$ and $j = 1, 2, \dots, q$.

In-training Chained Augmentation Mixing. Inspired by AugMix [12], we apply the training-time chained augmentation mixing technique to further make the data more diverse and also improve the generalization and robustness of the model. The number of augmentation operations in each augmentation chain is randomly set as one, two, or three. The parameters in the Beta distribution and the Dirichlet distribution are all set as 1. We create two sets of augmentation collections, namely AugMix-Soft and AugMix-Hard. Specifically, AugMix-Soft includes autocontrast, equalize, posterize and solarize, while AugMix-Hard has additional color, contrast, brightness, and sharpness augmentations. The overall expression of the synthetic training sample after the training-time augmentation mixing with N chains is

$$x_s^{AM} = m \cdot \Theta(x_f, x_b) + (1 - m) \cdot \sum_{i=1}^N (w_i \cdot H_i(\Theta(x_f, x_b))) \quad (2)$$

where m is a random convex coefficient sampled from a Beta distribution, w_i is also a random convex coefficient sampled from a Dirichlet distribution controlling the mixing weights of the augmentation chains. Both distribution functions

have the same coefficient value of 1. H_i denotes the integrated augmentation operations in the i^{th} augmentation chain.

3 Experiments

3.1 Datasets

Based on effortlessly collected source images and considering the contents in real surgery images, we apply a wide range of augmentation and blending operations to create abundant synthetic images for training. Only one background tissue image is adopted to generate our synthetic datasets. Specifically, for the case of 2 source images per instrument, we first organize the dataset Synthetic-A with 4000 synthetic images, and only one instrument exists in each synthetic image. Then we consider adding up additional 2000 synthetic images to build the dataset Synthetic-B where each image contains 2 distinct instruments. Moreover, we utilize one more source foreground image for each instrument and generate 2000 more synthetic images, among which 80% contain one instrument, and the remaining 20% contain 2 different instruments. This dataset with 8000 samples in total is named Synthetic-C.

To evaluate the quality of the generated surgical scene dataset, we conduct binary segmentation experiments with our synthetic datasets and the real EndoVis-2018 [1] dataset. We also evaluate on EndoVis-2017 [2] dataset to show that the model trained with our synthetic dataset also obtains good generalization ability to handle new domains with unseen instruments like the Vessel Sealer.

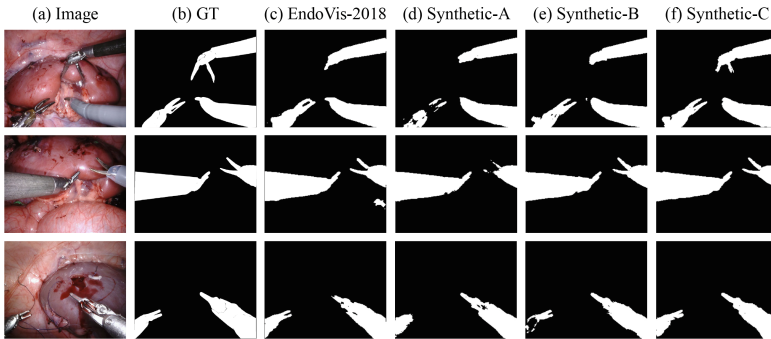


Fig. 2. Qualitative comparison of the binary segmentation results. (b) represents the ground truth. (c), (d), (e), and (f) show the results obtained from models trained with the EndoVis-2018 dataset and our three synthetic datasets.

3.2 Implementation Details

The classic state-of-the-art encoder-decoder network UNet [17] is used as our segmentation model backbone. We adopt a vanilla UNet architecture³ with Pytorch [16] library and train the model with NVIDIA RTX3090 GPU. The batch size of 64, the learning rate of 0.001, and the Adam optimizer are identically used for all experiments. The binary cross-entropy loss is adopted as the loss function. We use the Dice Similarity Coefficient (DSC) to evaluate the segmentation performance. The images are resized to 224×224 to save the training time. Besides, we refer to the implementation⁴ of AugMix [12] to apply training-time chained augmentation mixing.

3.3 Results and Discussion

We evaluate the quality and effectiveness of our generated dataset with the EndoVis-2018 [1] and EndoVis-2017 [2] datasets, with the latter one considered as an unseen target domain because it does not contribute to our synthetic dataset generation. The results in Table 1 indicate that our methods can complete the segmentation task with acceptable performance for both datasets. As shown in Fig. 2, the instrument masks predicted by our models only have minimal visual discrepancy from the ground truth. Considering our datasets only depend on a few trivially collected source images and get rid of gathering and annotating hundreds of real data samples, the result is promising and revolutionary for low-cost and efficient surgical instrument segmentation.

Table 1. Overall results of the binary surgical instrument segmentation in DSC (%) with the EndoVis-2018 dataset and our three synthetic datasets. AM is short for the training-time augmentation mixing. Best results of ours are shown in bold.

Train	Test on EndoVis-2018			Test on EndoVis-2017		
	AM-None	AM-Soft	AM-Hard	AM-None	AM-Soft	AM-Hard
EndoVis-2018	81.58	83.15	82.91	83.21	84.06	83.43
Synthetic-A	56.82	66.74	71.03	72.74	72.23	65.13
Synthetic-B	57.37	69.42	72.53	72.65	73.21	72.41
Synthetic-C	59.28	71.48	73.51	74.37	75.69	75.16

3.4 Ablation Studies

To show the efficacy of our training-time chained augmentation mixing, we first conduct experiments with a relevant data augmentation technique - ColorJitter,

³ <https://github.com/ternaus/robot-surgery-segmentation>.

⁴ <https://github.com/google-research/augmix>.

which randomly changes the brightness, contrast, and saturation of an image. Training with the Synthetic-C dataset, our augmentation strategy outperforms ColorJitter significantly with 5.33% and 4.29% of DSC gain on EndoVis-2018 and EndoVis-2017 datasets.

We then study the effectiveness of training with synthetic data in handling the class-incremental issue in the deployment domain. Compared with EndoVis-2018 [1] dataset, there are two novel instruments in EndoVis-2017 [2], namely the Vessel Sealer and the Grasping Retractor. Following our proposed framework in Fig. 1, we generate 2000 synthetic images for the novel instruments and combine them with EndoVis-2018 [1] for training. As indicated in the highlighted area of Fig. 3(a), the model manages to handle the class-incremental problem to recognize the Vessel Sealer, with only minimal effort of adding synthesized images. The overall performance on the test domain improves significantly, as shown in Fig. 3(b).

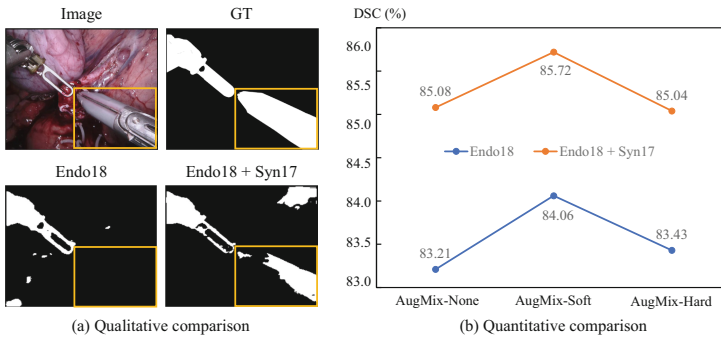


Fig. 3. Qualitative and quantitative results of the class-incremental case. The novel instrument Vessel Sealer is highlighted with the yellow rectangle in (a). The overall performance on EndoVis-2017 [2] gets greatly improved as shown in (b). (Color figure online)

While sufficient well-annotated datasets are not common in practice, a few high-quality data samples are normally feasible to acquire. We further investigate the effect of introducing a small portion of real images when training with synthetic data. We randomly fetch 10% and 20% of the EndoVis-2018 [1] dataset and combine it with our Synthetic-C dataset. The results in Table 2 indicate that only a small amount of real data could provide significant benefits. Compared with training with the real EndoVis-2018 [1] dataset, the models from the synthetic-real joint training scheme can efficiently achieve similar performance regarding adaptation and generalization.

Table 2. Results of synthetic-real joint training. Adding only a small portion of real data can greatly improve the segmentation performance.

Test	Train			
	Synthetic-C	Synthetic-C + 10% Endo18	Synthetic-C + 20% Endo18	Endo18
EndoVis-2018	73.51	80.65	82.45	82.91
EndoVis-2017	75.16	82.10	82.48	83.43

4 Conclusion

In this work, we reevaluate the surgical instrument segmentation and propose a cost-effective data-centric framework for synthetic dataset generation. Extensive experiments on two commonly seen real datasets demonstrate that our high-quality synthetic datasets are capable of surgical instrument segmentation with acceptable performance and generalization ability. Besides, we show that our method can handle domain shift and class incremental problems and greatly improve the performance when only a small amount of real data is available. Future work may be extended to more complicated instrument-wise segmentation and other medical applications. Besides, by considering more prior knowledge in practical surgical scenes, such as cautery smoke and instruments shadow, the quality of the synthetic dataset can be further improved.

Acknowledgements. This work was supported by the Shun Hing Institute of Advanced Engineering (SHIAE project BME-p1-21) at the Chinese University of Hong Kong (CUHK), Hong Kong Research Grants Council (RGC) Collaborative Research Fund (CRF C4026-21GF and CRF C4063-18G), (GRS)#3110167 and Shenzhen-Hong Kong-Macau Technology Research Programme (Type C 202108233000303).

References

1. Allan, M., et al.: 2018 robotic scene segmentation challenge (2020)
2. Allan, M., et al.: 2017 robotic instrument segmentation challenge (2019)
3. Cao, B., Zhang, H., Wang, N., Gao, X., Shen, D.: Auto-gan: self-supervised collaborative learning for medical image synthesis. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 10486–10493 (2020)
4. Colleoni, E., Edwards, P., Stoyanov, D.: Synthetic and real inputs for tool segmentation in robotic surgery. In: Martel, A.L., et al. (eds.) MICCAI 2020. LNCS, vol. 12263, pp. 700–710. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59716-0_67
5. Dobbs, R.W., Halgrimson, W.R., Talamini, S., Vigneswaran, H.T., Wilson, J.O., Crivellaro, S.: Single-port robotic surgery: the next generation of minimally invasive urology. *World J. Urol.* **38**(4), 897–905 (2020)
6. Domingos, P.: A few useful things to know about machine learning. *Commun. ACM* **55**(10), 78–87 (2012)

7. Eilertsen, G., Tsirikoglou, A., Lundström, C., Unger, J.: Ensembles of gans for synthetic training data generation (2021)
8. Garcia-Peraza-Herrera, L.C., Fidon, L., D’Ettorre, C., Stoyanov, D., Vercauteren, T., Ourselin, S.: Image compositing for segmentation of surgical tools without manual annotations. *IEEE Trans. Med. Imaging* **40**(5), 1450–1460 (2021)
9. Goodfellow, I., Bengio, Y., Courville, A.: *Deep Learning*. MIT press, Cambridge (2016)
10. Hamghalam, M., Lei, B., Wang, T.: High tissue contrast MRI synthesis using multi-stage attention-gan for segmentation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, pp. 4067–4074 (2020)
11. Han, C., et al.: Synthesizing diverse lung nodules wherever massively: 3d multi-conditional gan-based CT image augmentation for object detection. In: *2019 International Conference on 3D Vision (3DV)*, pp. 729–737. IEEE (2019)
12. Hendrycks, D., Mu, N., Cubuk, E.D., Zoph, B., Gilmer, J., Lakshminarayanan, B.: Augmix: a simple data processing method to improve robustness and uncertainty. *arXiv preprint arXiv:1912.02781* (2019)
13. Jung, A.B., et al.: imgaug. <https://github.com/aleju/imgaug>. Accessed 01 Feb 2020 (2020)
14. Kishore, A., Choe, T.E., Kwon, J., Park, M., Hao, P., Mittel, A.: Synthetic data generation using imitation training. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3078–3086 (2021)
15. Madan, S., et al.: When and how do cnns generalize to out-of-distribution category-viewpoint combinations? *arXiv preprint arXiv:2007.08032* (2020)
16. Paszke, A., et al.: Automatic differentiation in pytorch. In: *NIPS-W* (2017)
17. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
18. Shin, H.-C., et al.: Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: Gooya, A., Goksel, O., Oguz, I., Burgos, N. (eds.) *SASHIMI 2018*. LNCS, vol. 11037, pp. 1–11. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00536-8_1
19. Tremblay, J., et al.: Training deep networks with synthetic data: bridging the reality gap by domain randomization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 969–977 (2018)
20. Xu, M., Islam, M., Lim, C.M., Ren, H.: Class-incremental domain adaptation with smoothing and calibration for surgical report generation. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *MICCAI 2021*. LNCS, vol. 12904, pp. 269–278. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-87202-1_26