

Simulation-to-Real domain adaptation with teacher-student learning for endoscopic instrument segmentation

Manish Sahu · Anirban Mukhopadhyay ·
Stefan Zachow

Received: date / Accepted: date

Abstract

Purpose: Segmentation of surgical instruments in endoscopic videos is essential for automated surgical scene understanding and process modeling. However, relying on fully supervised deep learning for this task is challenging because manual annotation occupies valuable time of the clinical experts.

Methods: We introduce a teacher-student learning approach that learns jointly from annotated simulation data and unlabeled real data to tackle the erroneous learning problem of the current consistency-based unsupervised domain adaptation framework.

Results: Empirical results on three datasets highlight the effectiveness of the proposed framework over current approaches for the endoscopic instrument segmentation task. Additionally, we provide analysis of major factors affecting the performance on all datasets to highlight the strengths and failure modes of our approach.

Conclusion: We show that our proposed approach can successfully exploit the unlabeled real endoscopic video frames and improve generalization performance over pure simulation-based training and the previous state-of-the-art. This takes us one step closer to effective segmentation of surgical tools in the annotation scarce setting.

Keywords Surgical instrument segmentation · Self-supervision · Consistency learning · Self-ensembling · Unsupervised domain adaptation

Funded by the German Federal Ministry of Education and Research (BMBF) under the project COMPASS (grant no. - 16 SV 8019).

M. Sahu · S. Zachow
Zuse Institute Berlin (ZIB), Germany
E-mail: {sahu, zachow}@zib.de

A. Mukhopadhyay
Department of Computer Science, TU Darmstadt, Germany
E-mail: anirban.mukhopadhyay@gris.tu-darmstadt.de

1 Introduction

A faithful segmentation of surgical instruments in endoscopic videos is a crucial component of surgical scene understanding and realization of automation in computer- or robot-assisted intervention systems.¹ A majority of recent approaches address the problem of surgical instrument segmentation by training deep neural networks (DNNs) in a fully-supervised scheme. However, the applicability of such supervised approaches is restricted by the availability of a sufficiently large amount of real videos with clean annotations. The annotation process (especially pixel-wise) can be prohibitively expensive (see Figure 1) because it takes valuable time of medical experts.

An alternative direction to mitigate the dependency on annotated video sequences is to utilize synthetic data for the training of DNNs. Recent advances in graphics and simulation infrastructures have paved the way to automatically create a large number of photo-realistic simulated images with accurate pixel-level labels.^{1;2;3} However, the DNNs trained purely on simulated images do not generalize well on real endoscopic videos due to domain shift issue.^{4;5} We hypothesize that a DNN's bias towards recognizing textures rather than shapes⁶ results in a significant drop of performance when the DNNs are trained on simulation (rendered) data and applied to real environments. This is mainly because the heterogeneity of information within a real surgical scene is heavily influenced by factors such as lighting conditions, motion blur, blood, smoke, specular reflection, noise etc. However, simulation data only mimic shapes of instrument and patient-specific organs.⁷ To alleviate this performance generalization issue, domain adaptation approaches⁸ are utilized to address domain shift/mismatch.

When the annotations for the target domain are not available for learning, domain adaptation is applied in an unsupervised manner,⁹ where enabling the DNN to learn domain-invariant features is the primary goal of a learning algorithm. For instance, Pfeiffer et al.³ utilized an *image-to-image* translation approach, where the simulated images are translated into realistic looking ones by mapping image styles (texture, lighting) of the real data using a Cycle-GAN. In contrast, we argued for a shape-focused joint learning from simulated and real data in an end-to-end fashion and introduced a consistency-learning based approach *Endo-Sim2Real*¹⁰ to align the DNNs on both domains. We showed that similar performance can be

¹ EndoVis Sub-challenges - 2015, 2017, 2018, 2019 [<https://endovis.grand-challenge.org>]

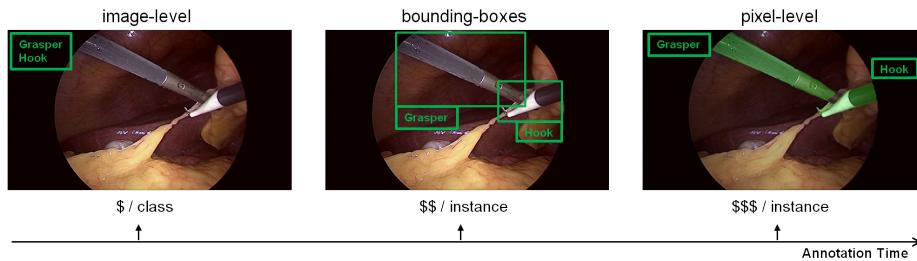


Fig. 1: Fully supervised deep learning is unrealistic for instrument segmentation due to a significantly high annotation effort.

obtained by employing a non-adversarial approach while improving the computational efficiency (with respect to training time). However, similar to perturbation-based consistency learning approaches,^{11;12} *Endo-Sim2Real* suffers from so-called confirmation bias.¹³ This is caused by noise accumulation or erroneous learning during the training stages, which may result in a degenerate solution.¹⁴

In this work, we formalise the consistency-based unsupervised domain adaptation framework to identify the confirmation bias problem of *Endo-Sim2Real* and propose a teacher-student learning paradigm to address this problem. Our proposed approach tackles the erroneous learning by improving the pseudo-label generation procedure for the unlabeled data and facilitate stable training of DNNs while maintaining computational efficiency. Through quantitative and qualitative analysis, we show that our proposed approach outperforms the previous *Endo-Sim2Real* approach across three data sets. Moreover, the proposed approach leads to a stable training without loosing computational efficiency.

Our contribution in this article is two-fold:

1. We introduce the teacher-student learning paradigm to the task of surgical instrument segmentation in endoscopic videos.
2. We provide a thorough quantitative and qualitative analysis to show the failure modes of our unsupervised domain adaptation approach. Our analysis leads to a better understanding of the strengths and challenges of consistency-based unsupervised learning with simulation-based supervision.

2 Related Work

Research on instrument segmentation for endoscopic procedures is dominated by supervision-based approaches ranging from full supervision¹⁵, semi/self supervision¹⁶, and weak supervision¹⁷ up to multi-task¹⁸ and multi-modal learning¹⁹. Some recent works also explored unsupervised approaches,^{20;21} however, for the sake of brevity, we will only focus on approaches that employ learning from simulation data for unsupervised domain adaptation.

Within the context of domain adaptation in surgical domains, Mahmood et. al.²² proposed an adversarial-based transformer network to translate a real image to a synthetic image such that a depth estimation model trained on synthetic images can be applied to the real image. On the other hand, Rau et. al.²³ proposed a conditional Generative Adversarial Network (GAN) based approach to estimate depth directly from real images. Other works have argued for translating synthetic images to photo-realistic images by using domain mapping via style transfer^{24;25}, for instance by using Cycle-GAN based *unpaired image-to-image* translation^{26;27} and utilize annotations from synthetic environment for deep learning tasks.

Pfeiffer et. al.³ proposed an unpaired *image-to-image* translation approach *I2I* that focuses on reducing the distribution difference between the source and the target domain by employing a Cycle-GAN based style transfer. Afterwards, a DNN is trained on the the translated images and its corresponding labels. On the other hand, *Endo-Sim2Real*¹⁰ utilizes similarity-based joint learning from both simulation and real data under the assumption that the shape of an instrument remains consistent across domains as well as under semantic preserving perturbations (like adding pixel-level noise or transformations).

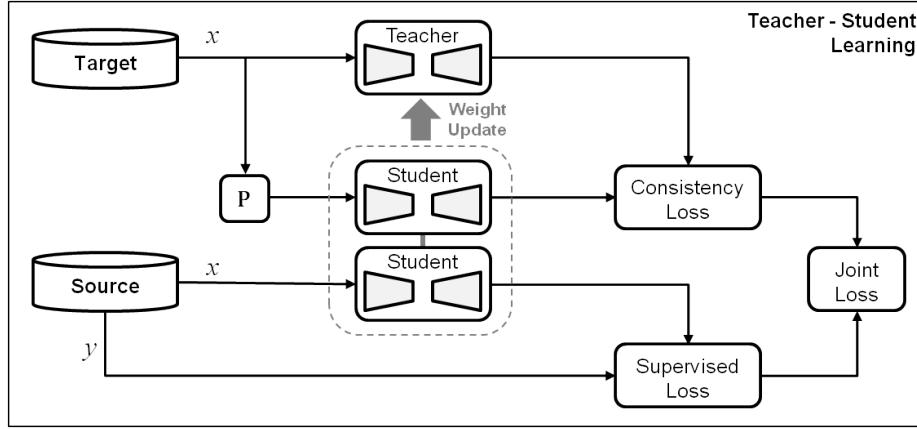


Fig. 2: Proposed teacher-student learning approach comprising supervised learning from simulated data as well as consistency learning from real unlabeled data.

This work is in line with *Endo-Sim2Real* and focuses on end-to-end learning for unsupervised domain adaptation (UDA). However, we formalise the consistency-based UDA to identify the confirmation bias problem and unstable training of *Endo-Sim2Real* approach and address it by employing a teacher-student paradigm. This facilitates stable training of the DNN and enhances its performance generalization capability.

3 Method

Our proposed teacher-student domain-adaptation approach (see Figure 2) aims to bridge the domain gap between source (simulated) and target (real) data by aligning a DNN model to both domains. Given:

- a source domain $D_s = (X_s, Y_s)$ associated with a feature space X_s and a label space Y_s and containing n_s labeled samples $\{(x_i^s, y_i^s)\}_{i=1}^{n_s}$ where, $x_i \in X_s$ and $y_i \in Y_s$ denote the i -th pair of image and label data respectively
- a target domain $D_t = (X_t)$ associated with a feature space $|X_t|$ and a label space $|Y_t|$ and containing n_t unlabeled samples $\{x_i^t\}_{i=1}^{n_t}$ where, $x_i \in X_t$ denote the i -th image of the unlabeled data

the goal of unsupervised domain adaptation is to learn a DNN model that generalizes on the target domain D_t . It is important to note that although the simulation and real endoscopic scene may appear similar, the label space between source- and target-domains generally differ (i.e. $Y_s \neq Y_t$), representing for example different organs or different instrument types. Since we are focusing on binary instrument segmentation, the label categories are twofold (i.e. $Y_s = Y_t = \{\text{"instrument"}, \text{"background"}\}$). For the sake of simplicity, we refer to the source domain D_s as labeled simulated domain D_L^{Sim} and to the target domain D_t as unlabeled real domain D_{UL}^{Real} .

Our proposed (and previous) approach learns by jointly minimizing the supervised loss L_{sl} for the labeled simulated data-pair as well as the consistency loss

L_{cl} for the unlabeled real data. A core component of the joint learning approach is unsupervised consistency learning, where a supervisory signal is generated by enforcing the DNN f_θ (parameterized with network weights θ) to produce a consistent output for an unlabeled input x and its perturbed form $\mathcal{P}(x)$.

$$\min_{\theta} \mathcal{L}_{sl} + \mathcal{L}_{cl} \underbrace{\left\{ f_\theta(x), f_\theta(\mathcal{P}(x)) \right\}}_{\tilde{y}} \quad (1)$$

Here in Equation 1, the DNN prediction \tilde{y} for unperturbed data x acts as a *pseudo-label* for perturbed data $\mathcal{P}(x)$ to guide the learning process. Therefore, the *Endo-Sim2Real* scenario can be interpreted as a *student-as-teacher* approach where the DNN acts as both a teacher that produces pseudo-labels and a student that learns from these labels. Since the DNN predictions may be incorrect or noisy during training,¹¹ this *student-as-teacher* approach leads to so-called the confirmation bias,¹³ which reinforces the student to overfit to the incorrect pseudo-labels generated by the teacher and prevents learning new information. This issue is especially prominent during early stages of the training, when the DNN still lacks the correct interpretation of the labels. If the unsupervised consistency loss (L_{cl}) outweighs the supervised loss (L_{sl}), the learning process is not effective and leads to a sub-optimal performance. Therefore, the consistency loss is typically employed with a temporal weighting function $w(t)$ such that the DNN learns prominently from the supervised loss during the initial stages of the learning and gradually shifts towards unsupervised consistency learning in the later stages.

Although the temporal ramp-up weighting function in *Endo-Sim2Real* helps to reduce the effect of the confirmation bias during joint learning, the DNN still learns directly from the incorrect pseudo-labels generated by the teacher.

In this work, we address this major drawback of the *Endo-Sim2Real* approach by improving the pseudo-label generation procedure of the unlabeled consistency learning. To this end, the teacher network is de-coupled from the student network and redefined ($f_\theta \rightarrow f_{\theta'}$) to generate reliable targets to enable the student to gradually learn meaningful information about the instrument shape. In order to avoid separate training of the teacher model, the same architecture is used for the teacher and its parameters are updated as a temporal average¹³ of the student network's weights.

$$\min_{\theta} \left\{ \underbrace{\sum_{(x_i, y_i) \in D_L^{Sim}} \mathcal{L}_{sl}(f_\theta(x_i), y_i)}_{supervised\ simulation} + w(t) * \underbrace{\sum_{x_i \in D_{UL}^{Real}} \mathcal{L}_{cl}(f_{\theta'}(x_i), f_\theta(\mathcal{P}(x_i)))}_{unsupervised\ real} \right\}$$

$$\dot{\theta}_t' = (\alpha \cdot \dot{\theta}_{t-1}' + (1 - \alpha) \cdot \theta_t) \quad (2)$$

At each training step t , the student f_θ is updated using *gradient-descent* while the teacher $f_{\theta'}$ is updated using student network weights, where the smoothing factor α controls the update rate of the teacher. A pseudo code of our proposed teacher-student learning approach is provided in Algorithm 1.

Algorithm 1 Teacher-student Algorithm for Domain Adaptation

Model: f_θ ▷ student with trainable parameters θ .
Model: $f_{\theta'}$ ▷ teacher with trainable parameters θ' .
Data: $D_L^{Sim}(x, y)$ ▷ labeled simulated samples
Data: $D_{UL}^{Real}(x)$ ▷ unlabeled real samples
Require: α ▷ update rate of teacher
Require: $w(t)$ ▷ temporal weight of consistency loss
Ensure: $\theta' \leftarrow \theta$ ▷ Initialize weights

```

1: loop JOINT LEARNING( $D_L^{Sim}, D_{UL}^{Real}$ ) ▷ ( $D_s \rightarrow D_L^{Sim}, D_t \rightarrow D_{UL}^{Real}$ )
2:   Supervised Loss
3:    $\{(x_i, y_i)\}_{i=1}^B \sim D_L^{Sim}(x, y)$  ▷ Sample mini-batch
4:    $\{(x_i^a, y_i^a)\}_{i=1}^B = \{\mathcal{A}(x_i, y_i)\}_{i=1}^B$  ▷ Augment batch
5:    $L_{sl} = \{f_\theta(x_i^a), y_i^a\}_{i=1}^B$  ▷ Supervised loss
6:   Consistency Loss
7:    $\{x_i\}_{i=1}^B \sim D_{UL}^{Real}(x)$  ▷ Sample mini-batch
8:    $\{x_i^p\}_{i=1}^B = \{\mathcal{P}(x_i)\}_{i=1}^B$  ▷ Perturb batch
9:    $\{\tilde{y}_i\}_{i=1}^B = \{f_{\theta'}(x_i^p)\}_{i=1}^B$  ▷ Pseudo Segmentation
10:   $L_{cl} = \{f_\theta(x_i), \tilde{y}_i\}_{i=1}^B$  ▷ Unsupervised loss
11:  Joint Loss
12:   $L = L_{sl} + w(t) \cdot L_{cl}$  ▷ Joint loss
13:   $g_\theta = \nabla_\theta L$  ▷ Compute gradients
14:   $\theta \leftarrow Update(\theta, g_\theta)$  ▷ Update student (gradient descent)
15:   $\theta' \leftarrow (\alpha \cdot \theta' + (1 - \alpha) \cdot \theta)$  ▷ Update teacher
16: end loop
    return  $\theta'$  ▷ Learned model

```

4 Experimental Setup**4.1 Data**

Simulation³ data contain $20K$ rendered images acquired via 3-D laparoscopic simulations from the CT scans of 10 patients. The images describe a rendered view of a laparoscopic scene with each tissue having a distinct texture and a presence of two conventional surgical instruments (grasper and hook) under a random placement of the camera (coupled with a light source).

Cholec¹⁰ data contain around $7K$ endoscopic video frames acquired from 15 videos of the Cholec80 dataset.²⁸ The images describe the laparoscopic cholecystectomy scene with seven conventional surgical instruments (grasper, hook, scissors, clipper, bipolar, irrigator and specimen bag). The data provide segmentations for each instrument type, however, the specimen bag is considered as a counterexample that is treated as background during evaluation, following the definition of an instrument in RobustMIS challenge.²⁹

EndoVis³⁰ data consist of 300 images from six different in-vivo 2D recordings of complete laparoscopic colorectal surgeries. The data provide binary segmentations of instruments for validation where images describe an endoscopic scene containing seven conventional instruments (including hook, traumatic grasper, ligasure, stapler, scissors and scalpel).¹⁵

Table 1: List of source (simulation) and target (real) datasets used during evaluation, where [videos (#) | empty frames (%)] reflects the number of videos and percentage of frames with no instrument respectively

Dataset	Training	Testing	Instruments
Simulation	20,000 (10 33 %)	n.a.	two
Cholec15	5,034 (10 12 %)	2,136 (5 13 %)	six + specimenbag
EndoVis	160 (4 0 %)	140 (6 0 %)	seven
RobustMIS	5,983 (16 17 %)	4,057 (14 20 %)	six + trocar
Stage 1		663 (2 15 %)	-
Stage 2		514 (2 13 %)	-
Stage 3		2880 (10 23 %)	-

RobustMIS²⁹ data consist of around 10K images acquired from 30 surgical procedures of three different types of colorectal surgery (10 rectal resection procedures, 10 proctocolectomy procedures and 10 procedures of sigmoid resection procedures). An instrument is defined as an elongated rigid object that is manipulated directly from outside the patient. Therefore, grasper, scalpel, clip applicator, hooks, stapling device, suction and even trocar is considered as an instrument while non-rigid tubes, bandages, compresses, needles, coagulation sponges, metal clips etc. are considered as counterexamples as they are indirectly manipulated from outside.²⁹ The data provide instance level segmentations for validation, which are performed in three different stages with an increasing domain gap between the training- and the test-data. Stage 1 contains video frames from 16 cases of the training data, stage 2 has video frames of two proctocolectomy and rectal surgeries each, and stage 3 has video frames from 10 sigmoid resection surgeries.

It is important to note that the domain gap increases not only in the three stages of testing in Robust-MIS dataset, but also from *Simulation* towards *Real* datasets (EndoVis < Cholec < Robust-MIS) as the definition of instrument (and/or counterexample) changes along with other factors.

4.2 Implementation

We have redesigned the implementation of the *Endo-Sim2Real* framework in view of a teacher-student approach. To ensure a direct and fair comparison, we employ the same *TerNaus11*³¹ as a backbone segmentation model. Also, we utilize the best performing perturbation scheme (i.e. applying one of the *pixel-intensity* perturbation² followed by one of the *pixel-corruption* perturbation³) and the loss function (i.e. *cross-entropy* and *jaccard*) of *Endo-Sim2Real* for evaluation. All simulated input images and labels are first pre-processed with a stochastically-varying circular outer mask to give them the appearance of real endoscopic images.

² *pixel-intensity*: random brightness and contrast shift, posterisation, solarisation, random gamma shift, random HSV color space shift, histogram equalization and contrast limited adaptive histogram equalization

³ *pixel-corruption*: gaussian noise, motion blurring, image compression, dropout, random fog simulation and image embossing

We use a batch size of 8 for 50 epochs and apply weight decay ($1e - 6$) as standard regularization. During consistency training, we use a time-dependent weighting function, where the weight of the unlabeled loss term is linearly increased over the training. The teacher model is updated with α (0.95) at each training step.

During evaluation of a dataset, we use an image-based dice score and average over all images to obtain a global dice metric for the dataset. For computation of the dice score, we exclude the cases where both the prediction and ground truth images are empty. However, we include cases with false positives for the empty images and set it to zero. So the dice score for empty ground-truth images (without any instrument) is either zero and considered in case of any false positives or undefined and not considered in case of correct prediction. Also, we report all the results as an average performance of three runs throughout our experiments.

5 Results and Discussion

This section provides a quantitative comparison with respect to the state-of-the-art approaches to demonstrate the effectiveness of our approach. Moreover, we perform a quantitative and qualitative analysis on three different datasets with varying degrees of the domain gap. This shows the strengths and weaknesses of our approach in order to better understand the challenges and provide valuable insights into addressing the remaining performance gap.

5.1 Comparison with *baseline* and *state-of-the-art*

In these experiments, we first highlight the performance of the two baselines: the *lower baseline* (supervised learning purely on simulated data) and the *upper baseline* (supervised learning purely on annotated real data) in Table 2. The substantial performance gap between the baselines indicates the domain gap between simulated and real data. Secondly, we compare our proposed teacher-student approach with other unsupervised domain adaptation approaches, i.e. the domain style transfer approach (*I2I*) and the plain consistency-based joint learning approach (*Endo-Sim2Real*) on the *Cholec* dataset. The empirical results show that *Endo-Sim2Real* works similar to *I2I*, while our proposed approach outperforms

Table 2: Quantitative comparison using DSC [mean (std)]. The two baselines: simulation only and real only means training only on the simulated data and training on the annotated real data respectively (i.e. no adaptation). The paired t-test for *Endo-Sim2Real* and our work results in $p\text{-value} << 0.01$.

Approach	EndoVis	Cholec	R-MIS (S 1)	R-MIS (S 2)	R-MIS (S 3)
Simulation only	.42 (.29)	.30 (.30)	.30 (.28)	.36 (.30)	.20 (.24)
I2I ³	n.a.	.68 (.30)	n.a.	n.a.	n.a.
Endo-Sim2Real ¹⁰	.76 (.17)	.68 (.31)	.57 (.33)	.60 (.32)	.51 (.35)
Teacher-Student	.80 (.16)	.72 (.30)	.61 (.32)	.65 (.31)	.58 (.34)
Real only	.93 (.06)	.86 (.24)	.82 (.27)	.83 (.25)	.78 (.30)
Empty frames	0 %	13 %	15 %	13 %	23 %

both of these approaches. Later, we evaluate our approach on two additional datasets and show that it consistently outperforms *Endo-Sim2Real*. These experiments demonstrate that the generalization performance of the DNN can be enhanced by employing unsupervised consistency learning on unlabeled data. Finally, the performance gap with the upper baseline calls for identification of the issues needed to bridge the remaining domain gap.

5.2 Analysis on EndoVis

Among the three datasets, our proposed approach performs best for EndoVis as shown in Table 2. A visual analysis of the low performing cases in Figure 3 highlights factors such as false detection on specular reflection, under-segmentation for small instruments, tool-tissue interaction and partially occluded instruments. These factors can in part be addressed by utilizing the temporal information of video frames.³²

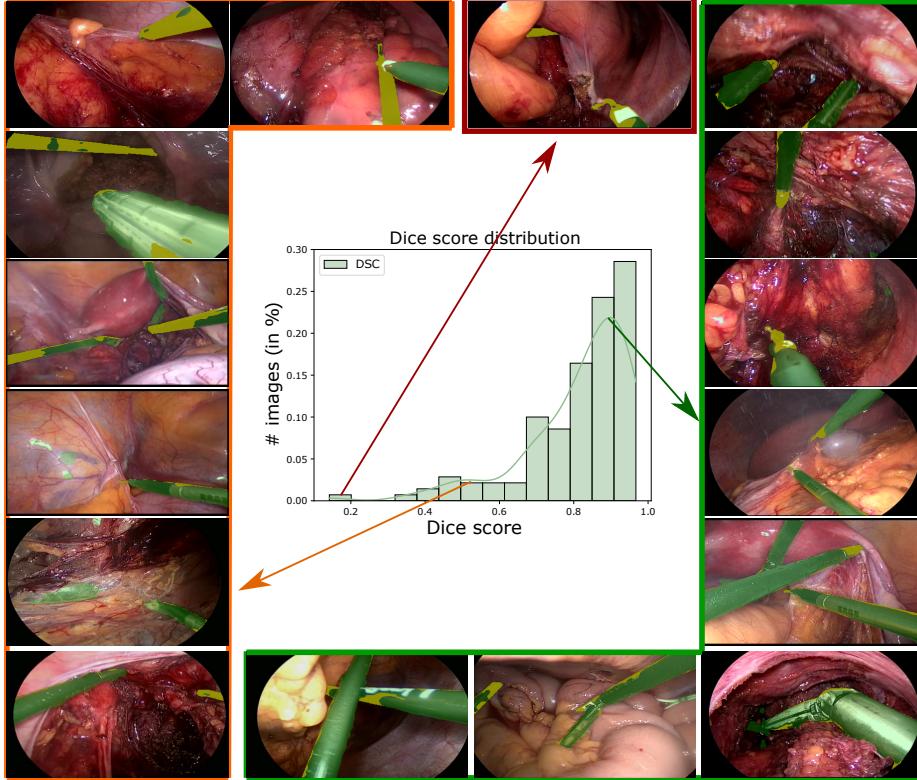


Fig. 3: Qualitative analysis on EndoVis dataset. The green color in the images represents the network predictions while the yellow color represents under-segmentation.

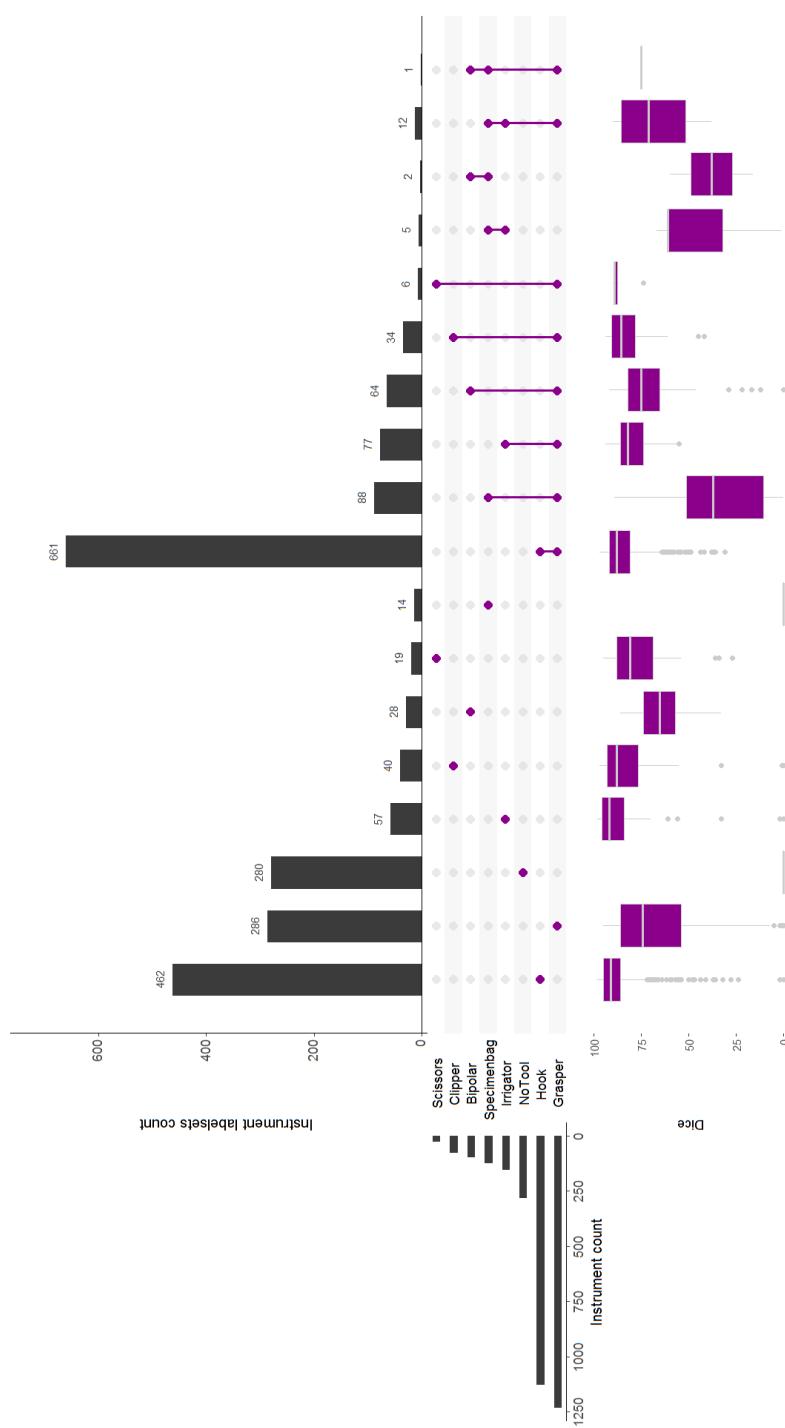


Fig. 4: Visualization of the relation between tool co-occurrence and segmentation quality for the Cholec dataset. Please note that the dice score is zero for no tool cases and specimen bag as it is treated as background.

5.3 Analysis on Cholec

We performed an extensive performance analysis of our proposed approach on the Cholec dataset as instrument-specific labels are available for it (in comparison to EndoVis). To understand the distinctive performance aspects for the Cholec dataset, we compare the segmentation performance across different instrument co-occurrence in Figure 4. A similar range of dice scores highlights that the performance of our approach is less impacted by the presence of multiple tool combinations in an endoscopic image. However, it also clearly shows that the segmentation performance of our approach drops when the specimen bag and its related co-occurrences are present (as seen in the respective box plots in Figure 4). A visual analysis highlights false detection on the reflective surface of the specimen bag.

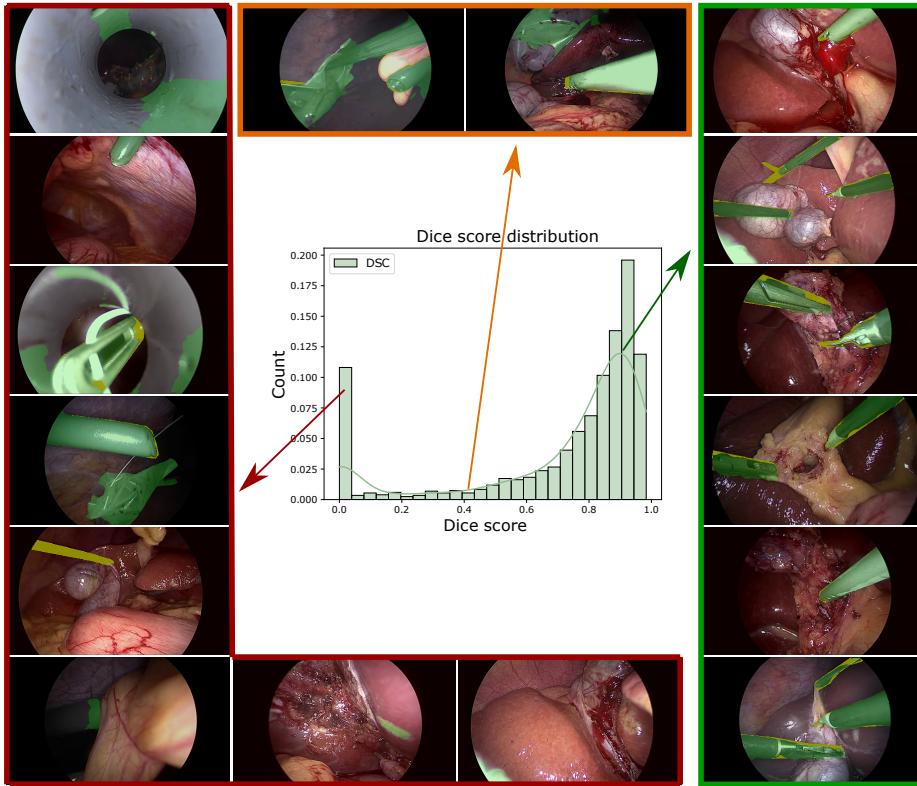


Fig. 5: Qualitative analysis on the Cholec dataset. The green color in the images represents the network predictions while the yellow color represents under-segmentation.

Apart from the previously analyzed performance degrading factors in the EndoVis dataset, other major factors affecting the performance are as follows:

- * **Out of distribution cases** such as a non-conventional tool-shape-like instrument: specimen bag (see box-plots for labelsets with specimen bag in Figure 4).

- * **False detection for scenarios** such as an endoscopic view within the trocar, instrument(s) near the image border or under-segmentation for small instruments.
- * **Artifact cases** such as specular reflection. The impact of other artefacts such as blood, smoke or motion blur is lower.

Although our proposed approach struggles to tackle these artefacts and out of distribution cases, addressing these performance degrading factors is itself an open research problem.³³

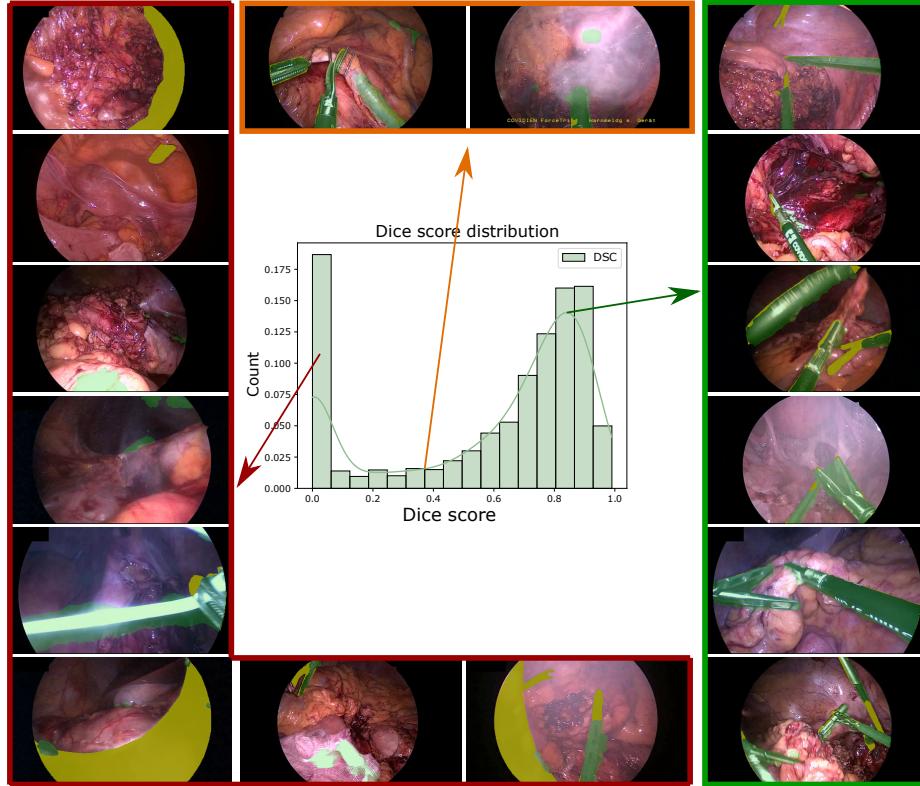


Fig. 6: Qualitative analysis on the RobustMIS dataset. The green color in the images represents the network predictions while the yellow color represents under-segmentation.

5.4 Analysis on RobusMIS

We analyzed the performance of our approach on images with a different number of instruments in the RobustMIS dataset. We found that the performance is not significantly affected by the presence of multiple tools (see Table 3). A low

Table 3: Quantitative results for multi-instrument presence in RobustMIS dataset using DSC [mean (std)]

No.	1	2	3	4	5	6	7
#	1802	1173	229	31	3	-	1
DSC	.61 (.34)	.72 (.22)	.73 (.18)	.76 (.11)	.76 (.13)	-	.77 (.00)

performance for a single visible instrument is attributed to small, stand-alone instruments across image boundary. Apart from the factors in the Cholec dataset, other real-world performance degrading factors in RobustMIS include: presence of other out-of-distribution cases such as non-rigid tubes, bandages, needles etc.; presence of corner cases such as trocar-views and specular reflections producing a instrument-shape-like appearance. These failure cases highlight a drawback of our approach, which works under the assumption that the shape of the instrument remains consistent between the domains. Therefore, our approach may not be able to produce faithful predictions in case instruments with different shapes are encountered in the real domain (compared to instruments in simulation) or counterexamples with instrument like appearance.

5.5 Impact of empty ground-truth frames

The performance of our teacher-student approach is negatively affected by the video frames that do not contain instruments. This is because the dice score is assigned to zero when the network predicts false positives (as seen in Figure 5 and 6) in instrument-free video frames. A direct relation of this effect can be seen in Table 2 where the dice score across the datasets decreases as the number of empty frames increases (in %) from EndoVis to RobustMIS. It suggests that utilizing false detection techniques in the current framework can help in enhancing the generalization capabilities.

6 Conclusion

We introduce teacher-student learning to address the confirmation bias issue of the *EndoSim2Real* consistency learning. This enables us to tackle the challenging problem of the domain shift between synthetic and real images for surgical tool segmentation in endoscopic videos. Our proposed approach enforces the teacher model to generate reliable targets to facilitate stable student learning. Since the teacher is a moving average model of the student, the extension does not add computational complexity to the current approach.

We show that the proposed teacher-student learning approach generalizes across three different datasets for the instrument segmentation task and consistently outperforms the previous state-of-the-art. For a majority of images (see high peak in Figure 3, 5 and 6), the segmentation predictions are usually correct with small variations across the instrument boundary. Moreover, a thorough analysis of the results highlight interpretable failure modes of simulation-to-real deep learning as the domain gap widens progressively.

Considering the strengths and weaknesses of our teacher-student enabled consistency-based unsupervised domain adaptation approach, the framework admits multiple straightforward extensions:

- * Implementing techniques to suppress false detection for empty frames, instruments near the image border and specular reflections, for instance by utilizing temporal information³² of video frames.
- * Improving physical properties of simulation to capture instrument-tissue interaction, considering the variations in predictions across instrument boundaries.
- * Extension towards other domain-adaptation-tasks, for instance depth estimation^{22;23} or instrument pose estimation^{34;35} by exploiting depth maps from the simulated virtual environments.
- * Extension towards semi-supervised domain adaptation or real-to-real unsupervised domain adaptation by joint learning from labeled and unlabeled data.
- * Utilizing this approach on top of self-supervised or adversarial domain mapping approaches such as *I2I*.

The heavy reliance of current approaches on manual annotation and the harsh reality of surgeons sparing time for the annotation process propels simulation-to-real domain adaptation as the obvious problem to address in surgical data science. The proposed approach ushers annotation-efficient surgical data science for the operating room of the future.

Author Statement

Conflict of interest: The authors state no conflict of interest.

Informed consent: This study contains patient data from a publicly available dataset.

Ethical approval: This article does not contain any studies with human participants or animals performed by any of the authors.

References

1. Hoffman J, Wang D, Yu F, Darrell T (2016) Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. arXiv preprint arXiv:161202649 2
2. Zhang Y, David P, Foroosh H, Gong B (2019) A curriculum domain adaptation approach to the semantic segmentation of urban scenes. IEEE transactions on pattern analysis and machine intelligence 2
3. Pfeiffer M, et al (2019) 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, Springer, pp 119–127 2, 3, 6, 8
4. Torralba A, Efros AA (2011) Unbiased look at dataset bias. In: CVPR 2011, IEEE, pp 1521–1528 2
5. Vercauteren T, Unberath M, Padoy N, Navab N (2020) Cai4cai: The rise of contextual artificial intelligence in computer-assisted interventions. Proceedings of the IEEE 108(1):198–214, DOI 10.1109/JPROC.2019.2946993 2

6. Baker N, Lu H, Erlikhman G, Kellman PJ (2018) Deep convolutional networks do not classify based on global object shape. *PLoS computational biology* 14(12):e1006613 2
7. Engelhardt S, Sharan L, Karck M, De Simone R, Wolf I (2019) Cross-domain conditional generative adversarial networks for stereoscopic hyperrealism in surgical training. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 155–163 2
8. Wang M, Deng W (2018) Deep visual domain adaptation: A survey. *Neurocomputing* 312:135–153 2
9. Wilson G, Cook DJ (2020) A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)* 11(5):1–46 2
10. Sahu M, Strömsdörfer R, Mukhopadhyay A, Zachow S (2020) Endo-sim2real: Consistency learning-based domain adaptation for instrument segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 784–794 2, 3, 6, 8
11. Laine S, Aila T (2017) Temporal ensembling for semi-supervised learning. In: International Conference on Learning Representations 3, 5
12. French G, Mackiewicz M, Fisher M (2018) Self-ensembling for visual domain adaptation. In: International Conference on Learning Representations 3
13. Tarvainen A, Valpola H (2017) Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In: Advances in neural information processing systems, pp 1195–1204 3, 5
14. Chapelle O, Scholkopf B, Zien A (2009) Semi-supervised learning. *IEEE Transactions on Neural Networks* 20(3):542–542 3
15. Bodenstedt S, Allan M, Agustinos A, Du X, Garcia-Peraza-Herrera L, Kenngott H, Kurmann T, Müller-Stich B, Ourselin S, Pakhomov D, et al (2018) Comparative evaluation of instrument segmentation and tracking methods in minimally invasive surgery. *arXiv preprint arXiv:180502475* 3, 6
16. Ross T, Zimmerer D, Vemuri A, Isensee F, Wiesenfarth M, Bodenstedt S, Both F, Kessler P, Wagner M, Müller B, et al (2018) Exploiting the potential of unlabeled endoscopic video data with self-supervised learning. *International journal of computer assisted radiology and surgery* 13(6):925–933 3
17. Fuentes-Hurtado F, Kadkhodamohammadi A, Flouty E, Barbarisi S, Luengo I, Stoyanov D (2019) Easylabels: weak labels for scene segmentation in laparoscopic videos. *International journal of computer assisted radiology and surgery* pp 1–11 3
18. Laina I, Rieke N, Rupprecht C, Vizcaíno JP, Eslami A, Tombari F, Navab N (2017) Concurrent segmentation and localization for tracking of surgical instruments. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 664–672 3
19. Jin Y, Cheng K, Dou Q, Heng PA (2019) Incorporating temporal prior from motion flow for instrument segmentation in minimally invasive surgery video. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 440–448 3
20. Colleoni E, Edwards P, Stoyanov D (2020) Synthetic and real inputs for tool segmentation in robotic surgery. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 700–710 3

21. Liu D, Wei Y, Jiang T, Wang Y, Miao R, Shan F, Li Z (2020) Unsupervised surgical instrument segmentation via anchor generation and semantic diffusion. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 657–667 3
22. Mahmood F, Chen R, Durr NJ (2018) Unsupervised reverse domain adaptation for synthetic medical images via adversarial training. *IEEE transactions on medical imaging* 37(12):2572–2581 3, 14
23. Rau A, Edwards PE, Ahmad OF, Riordan P, Janatka M, Lovat LB, Stoyanov D (2019) Implicit domain adaptation with conditional generative adversarial networks for depth prediction in endoscopy. *International journal of computer assisted radiology and surgery* 14(7):1167–1176 3, 14
24. Luengo I, Flouty E, Giataganas P, Wisanuvej P, Nehme J, Stoyanov D (2018) Surreal: Enhancing surgical simulation realism using style transfer. In: British Machine Vision Conference 2018, BMVC 2018, BMVA, pp 1–12 3
25. Marzullo A, Moccia S, Catellani M, Calimeri F, De Momi E (2020) Towards realistic laparoscopic image generation using image-domain translation. *Computer Methods and Programs in Biomedicine* p 105834 3
26. Engelhardt S, De Simone R, Full PM, Karck M, Wolf I (2018) Improving surgical training phantoms by hyperrealism: deep unpaired image-to-image translation from real surgeries. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 747–755 3
27. Oda M, Tanaka K, Takabatake H, Mori M, Natori H, Mori K (2019) Realistic endoscopic image generation method using virtual-to-real image-domain translation. *Healthcare Technology Letters* 6(6):214–219 3
28. Twinanda AP, Shehata S, Mutter D, Marescaux J, De Mathelin M, Padov N (2016) Endonet: A deep architecture for recognition tasks on laparoscopic videos. *IEEE transactions on medical imaging* 36(1):86–97 6
29. Ross T, Reinke A, Full PM, Wagner M, Kenngott H, Apitz M, Hempe H, Filimon DM, Scholz P, Tran TN, et al (2020) Robust medical instrument segmentation challenge 2019. arXiv preprint arXiv:200310299 6, 7
30. (2015) Endovis sub-challenge: Instrument segmentation and tracking. <https://endovissub-instrument.grand-challenge.org/>, accessed: 2020-10-28 6
31. Shvets AA, Rakhlis A, Kalinin AA, Iglovikov VI (2018) Automatic instrument segmentation in robot-assisted surgery using deep learning. In: 2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA), IEEE, pp 624–628 7
32. González C, Bravo-Sánchez L, Arbelaez P (2020) Isinet: An instance-based approach for surgical instrument segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, pp 595–605 9, 14
33. Ali S, Zhou F, Braden B, Bailey A, Yang S, Cheng G, Zhang P, Li X, Kayser M, Soberanis-Mukul RD, et al (2020) An objective comparison of detection and segmentation algorithms for artefacts in clinical endoscopy. *Scientific reports* 10(1):1–15 12
34. Du X, Kurmann T, Chang PL, Allan M, Ourselin S, Sznitman R, Kelly JD, Stoyanov D (2018) Articulated multi-instrument 2-d pose estimation using fully convolutional networks. *IEEE transactions on medical imaging* 37(5):1276–1287 14

35. Allan M, Ourselin S, Hawkes DJ, Kelly JD, Stoyanov D (2018) 3-d pose estimation of articulated instruments in robotic minimally invasive surgery. *IEEE transactions on medical imaging* 37(5):1204–1213 14