



One model to use them all: training a segmentation model with complementary datasets

Alexander C. Jenke^{1,2,3,4} · Sebastian Bodenstedt^{1,2,3,4,6} · Fiona R. Kolbinger^{5,6,7} · Marius Distler^{5,6} · Jürgen Weitz^{5,6} · Stefanie Speidel^{1,2,3,4,6}

Received: 28 February 2024 / Accepted: 8 April 2024 / Published online: 27 April 2024
© The Author(s) 2024

Abstract

Purpose Understanding surgical scenes is crucial for computer-assisted surgery systems to provide intelligent assistance functionality. One way of achieving this is via scene segmentation using machine learning (ML). However, such ML models require large amounts of annotated training data, containing examples of all relevant object classes, which are rarely available. In this work, we propose a method to combine multiple partially annotated datasets, providing complementary annotations, into one model, enabling better scene segmentation and the use of multiple readily available datasets.

Methods Our method aims to combine available data with complementary labels by leveraging mutual exclusive properties to maximize information. Specifically, we propose to use positive annotations of other classes as negative samples and to exclude background pixels of these binary annotations, as we cannot tell if a positive prediction by the model is correct.

Results We evaluate our method by training a DeepLabV3 model on the publicly available Dresden Surgical Anatomy Dataset, which provides multiple subsets of binary segmented anatomical structures. Our approach successfully combines 6 classes into one model, significantly increasing the overall Dice Score by 4.4% compared to an ensemble of models trained on the classes individually. By including information on multiple classes, we were able to reduce the confusion between classes, e.g. a 24% drop for stomach and colon.

Conclusion By leveraging multiple datasets and applying mutual exclusion constraints, we developed a method that improves surgical scene segmentation performance without the need for fully annotated datasets. Our results demonstrate the feasibility of training a model on multiple complementary datasets. This paves the way for future work further alleviating the need for one specialized large, fully segmented dataset but instead the use of already existing datasets.

Keywords Full scene segmentation · Multi-class segmentation · Surgical scene understanding · Dataset availability · Surgical data science · Computer-assisted surgery

✉ Alexander C. Jenke
alexander.jenke@nct-dresden.de

¹ Department of Translational Surgical Oncology, National Center for Tumor Diseases (NCT/UCC) Dresden, Fetscherstraße 74, Dresden, Germany

² German Cancer Research Center (DKFZ), Heidelberg, Germany

³ Faculty of Medicine and University Hospital Carl Gustav Carus, Technical University Dresden, Dresden, Germany

⁴ Helmholtz-Zentrum Dresden-Rossendorf (HZDR), Dresden, Germany

⁵ Department of Visceral, Thoracic and Vascular Surgery, University Hospital and Faculty of Medicine Carl Gustav Carus, Technical University Dresden, Fetscherstraße 74, Dresden, Saxony, Germany

Introduction

The understanding of the visible surgical scene is key for computer-assisted surgery (CAS) systems to understand the current situation and provide adapted assistance functions. One approach of achieving this is via the use of full scene semantic segmentation models, which are able to classify every visible part of the scene. These models provide the basis

⁶ Centre for Tactile Internet with Human-in-the-Loop (CeTI), Technical University Dresden, CeTI Exzellenz-Cluster, Dresden, Saxony, Germany

⁷ Weldon School of Biomedical Engineering, Purdue University, West Lafayette, IN, USA

for recognizing the current situation or actions and enable useful assistance functions.

In recent work, progress has been made in improving surgical semantic scene segmentation by the use of temporal context [1], stereo vision [2], simulated data [3], or weak label annotations [4]. Further steps towards solving this task in the surgical setting have been taken by the Robotic Scene Segmentation Challenge [5] and the HeiSurF Challenge [6]. Both EndoVis¹ sub-challenges provided annotations of 11 and 21 classes, respectively, including surgical tools and human anatomy, and challenged participants to semantically segment all of them.

Nevertheless, a major bottleneck for clinical translation of surgical data science (SDS) applications remains the availability of such datasets [7], due to the high amount of time required for experts, to create such segmentation annotations. This is even more challenging for full scene semantic segmentation, which requires a pixel-wise annotation of multiple classes for the complete frame.

This issue was sidestepped in the recently published Dresden Surgical Anatomy Dataset (DSAD) [8] by simply providing binary segmentations. The dataset contains 11 classes of anatomical structures, split into multiple subsets providing one class each. This way, the authors were able to publish over 13,000 expert-approved semantic segmentations.

At present, there are a handful of datasets providing different annotations relevant to the field of SDS [7]. These datasets can differ in the granularity of classes and the classes annotated in general, depending on the protocol used during annotation. Recent works have presented different methods that rely only on partial annotations for training a CT segmentation model, improving the usage of existing knowledge [9–12]. The used approaches range from dataset specific backbones and pseudolabel generation [12], over merging unlabelled classes with the background class and adding a mutual exclusion constraint [9], to simply masking unlabelled classes during loss calculation [10].

We assume that learning to segment multiple classes causes a single model to develop a better understanding of the scene, leading to better segmentation performance. We therefore propose the usage of information from mutual exclusivity, which can be applied on top of any state-of-the-art approach. In this work we apply it in addition to the masking during loss calculation [10]. We demonstrate the feasibility of training a model for multi-class organ segmentation of laparoscopic surgery images on multiple complementary datasets, thereby overcoming the data bottleneck challenge in SDS. The code and models are publicly available on https://gitlab.com/nct_tso_public/dsad-segmentation/.

¹ <http://endovis.org>.

Methods

In this section, we introduce our proposed method and define an upper and lower baseline to compare the performance of the proposed method to. A visualization of the architectures is shown in Fig. 1. Further, we provide a definition of how the average dice scores are calculated in this work.

Baselines

A naive approach to combine multiple datasets is to train one model per dataset and subsequently building an ensemble to merge their predictions into one final multi-class prediction. In case of one binary segmenting model per class, this kind of ensemble prediction can be achieved by applying the argmax over the different sigmoid outputs per pixel. This is followed by a threshold to determine whether the most likely class is predicted positive; otherwise, the background class is assigned. In this work, this approach will be used as the lower baseline and is referred to as ensemble (EN).

Alternatively, a single model can be trained on a single fully labelled dataset. As this requires a large fully labelled dataset, which is hard to obtain but provides the maximum amount of information, this approach will be used as the upper baseline in this work and is referred to as fully supervised (FS).

Implication-based labelling

To overcome the ambiguity inherent to merging multiple models and lost information between classes, while sidestepping the need for a single large fully annotated dataset, we propose to combine the classes provided by multiple datasets into one model. This reduces computational costs during inference as only one output from one model is required and lowers the requirements we pose against the dataset, as not all classes need to be included in every dataset. In addition, the model could benefit from the shared knowledge among the different classes.

The model takes single images of the datasets as input and outputs a class probability for every pixel and every class. For every pixel, the final class is selected based on the highest predicted class probability. The outputs are normalized by a sigmoid, and therefore each value is independent of the outputs of other classes. On one hand, this setup does not prevent the model from predicting multiple classes per pixel; instead, this has to be enforced by the loss during training. On the other hand, this setup allows us to access the predicted probability of every class separately, which is required for dealing with the problem of training from incomplete knowledge in the datasets, as each dataset provides only information on its contained class and not the classes introduced by other datasets.

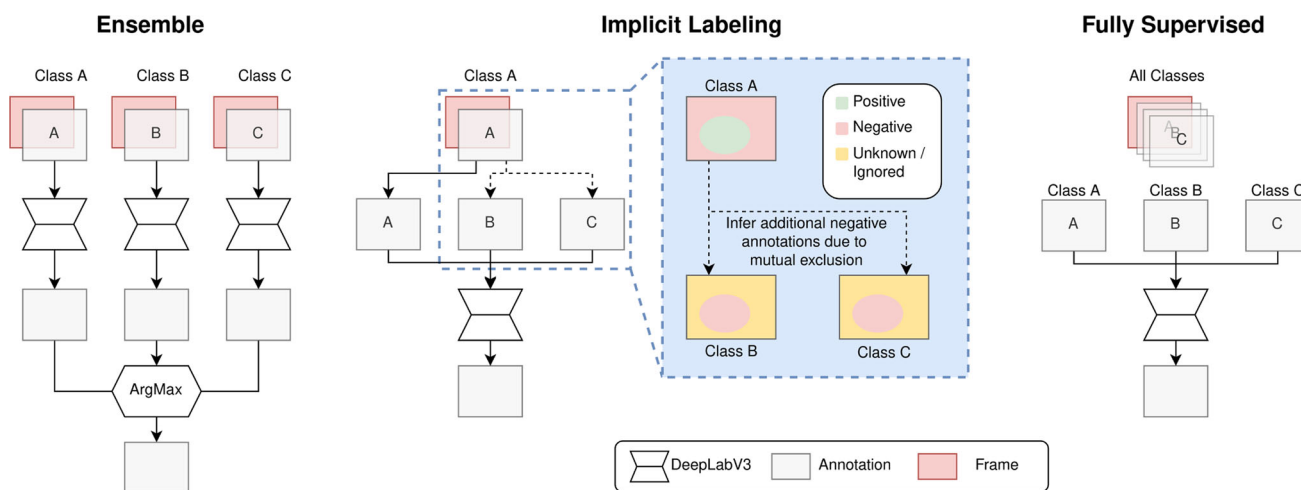


Fig. 1 A flowchart describing the flow of annotation data through the three architectures used in this work. The proposed implicit label generation is shown in the blue box

Assuming every pixel in the target semantic scene segmentation problem is exactly part of one class (mutually exclusivity of classes), this can be used to maximize the information provided by a partially annotated dataset by applying the following rules:

1. A positive annotation of one class implies a negative annotation of all other classes.
2. A negative annotation of a class (in a binary annotation the background) provides no information if this region contains other classes. As no implication to other classes can be made, the annotation stays *unknown*.

The application of these rules is visualized in the blue box of Fig. 1.

In this work, we use these implication rules to obtain additional negative samples from datasets not containing the examined class. Cases in which the annotation stays *unknown* are excluded from the loss calculation by masking, as no decision on the correctness of the prediction can be made. This method is called implicit labelling (IL) in the following.

The loss per class c can be formulated as:

$$L_c = \frac{1}{B} \frac{1}{\hat{P}} \sum_{b,p} BCE(\hat{y}_{b,p}^{(c)}, y_{b,p}^{(c)}) * \lambda_{b,p,c} \tag{1}$$

$$\lambda_{b,p,c} = \begin{cases} 1, & \text{if } b \text{ is annotated for } c \text{ or } y_{b,p}^{(\hat{c})} = 1, \hat{c} \in C \setminus c \\ 0, & \text{else} \end{cases} \tag{2}$$

where B is the number of images, $b \in [1, B]$, in the batch, and \hat{P} is the number of Pixels, $p \in P$, where $\lambda = 1$, $\hat{y}_{b,p}^{(c)}$ is the prediction of class c for pixel p of image b and $y_{b,p}^{(c)}$ the respective ground truth. A pixel p is included in the loss calculation if either the image b is annotated for class c or if $y_{b,p}^{(\hat{c})} = 1$ for any other class $\hat{c} \in C \setminus c$; otherwise, the pixel is

excluded by $\lambda_{b,p,c}$ being set to zero. In the second case the ground truth for class c is false, $y_{b,p}^{(c)} = 0$, due to the mutual exclusivity of classes.

Metrics

In this work, the model performance is evaluated using the dice score [13]. For the average dice score per class, the dice is calculated per image and averaged subsequently over all images. If a class does not occur in either the target or the prediction, the F1 is not defined; we therefore set the score to one in those cases as the model performed as expected. The average dice score over all classes is calculated by averaging all average dice scores per class.

Statistical significance is calculated using a two-sided Wilcoxon signed-rank test on the image-wise and class-wise dice scores of our approach against the lower baseline. The significance is calculated per class. The significance of the mean dice score is calculated by evaluating all image-wise and class-wise dice scores over all classes.

Evaluation

Dataset

We evaluate our proposed approach using the publicly available DSAD dataset [8]. This dataset consists of 13,195 laparoscopic images split into 11 subsets with a minimum of 1000 frames from a minimum of 20 surgeries each. In each subset, binary segmentations for one of the 11 classes (abdominal wall, colon, inferior mesenteric artery, intestinal veins, liver, pancreas, small intestine, spleen, stomach, ureter, vesicular glands) are provided. For the stomach subset, addi-

tional masks are available, annotating six of the remaining ten classes (abdominal wall, colon, liver, pancreas, small intestine, spleen) visible in this subset, resulting in one multi-class subset. This work uses the proposed split [14] into training, validation, and test set. To be able to compare to the multi-class subset, only the binary subsets of the contained classes are used, with the exception of the spleen. The spleen class was excluded in this work due to the lack of positive examples in the validation and test split in the multi-class subset. To examine the ability of IL and EN to join datasets, they were trained by splitting the multi-class subset into six binary sets, one for each class. The methods trained on these subsets are comparable to the FS approach trained on the combined annotations of the multi-class subset. As the binary subsets have no overlapping classes, we interpret them as separate datasets in this work, which we aim to join into one model. All classes fulfil the required mutual exclusiveness.

Trials

To validate our approach, we conducted five trials to compare our proposed implicit labelling approach to the ensemble approach and the fully supervised approach. In all experiments, the DeepLabV3 architecture with a ResNet50 backbone [15] is used. Models are initialized using the default PyTorch pretraining on COCO [16]. All models were trained using PyTorch [17] v1.12 on Ubuntu 18.04 and Nvidia Tesla V100 GPUs; evaluation was done on Ubuntu 20.04 with Nvidia RTX A5000 GPUs. The images were downscaled to a size of 640x512 pixels for memory and time reasons. For all trials an initial learning rate of 3×10^{-4} was used with a scheduler reducing the learn rate by a gamma of 0.9 every 10 epochs. Further, all models were trained for 100 epochs using an AdamW optimizer with weight decay of 0.1 and a cross-entropy loss with a *positive weight* factor for balancing positive and negative pixels. The best model per trial was selected according to the dice score on the validation set.

The ensemble approach evaluates an ensemble of six models, each trained on one of the classes, serving as the lower baseline. The prediction of the ensemble was obtained by selecting the class with the highest value over all models per pixel. The positive weight for the background was set to 1 and the remaining classes to the negative to positive pixel ratio per class. This approach was used in two trials, once trained on the binary subset and once on the binary sets extracted from the multi-class subset.

The fully supervised approach examines the fully supervised method, serving as the upper baseline. For this, a single DeepLabV3 was trained on the multi-class subset. The positive weight was calculated by the share of positive pixels of the class, negating the share, and normalizing all negated shares with the softmax function. As this approach requires

a fully annotated dataset, this method is only trained once on the multi-class subset.

In the implicit labelling approach, our proposed method was used to train a single DeepLabV3 on the six binary subsets, which represent the same classes that are available in the multi-class subset. The positive weight was set to the negative-to-positive pixel ratio while including the positive pixels of other classes as negatives, as described before. The output of the loss function was masked before aggregation to ignore pixels as required by our approach. This approach was used in two trials, once trained on the binary subset, and once on the binary sets extracted from the multi-class subset.

Results

The trials were evaluated twice on the up-to-now unseen test split of the dataset, once using the binary subset for each class and once using the multi-class subset. As the multi-class subset is based on the stomach subset, the frames and, therefore, the results for this class are identical.

The results on the binary testset are shown in Table 1. The overall best performance is reached by our proposed IL approach on the binary trainset with an average dice score of 72% over all classes outperforming EN by 4%. On the binary trainset, all scores of the IL approach are either significantly higher or not significantly lower than the EN approach. For the trials trained on the multi-class subset the upper baseline, FS, reaches the highest score of 46%. For the highest score per class distributes among the FS and IL approach. On the multi-class trainset, all scores of the IL approach are significantly above the EN approach, except for the stomach. The performance of all methods is significantly lower for Pancreas and Liver on the multi-class trainset compared to the binary trainset, while IL maintains the highest score.

The results on the multi-class testset are shown in Table 2. The overall best performance is reached by the FS approach with an average dice score of 78%. On the binary trainset, IL outperforms EN by 7%. Especially in the classes Colon, Pancreas, Small Intestine, and Stomach IL performs significantly better than EN, while EN is better for the Liver. On the multi-class testset FS has the highest score on all classes. The second-highest class is distributed among IL and EN. The average performance of IL outperforms EN by 4%.

Figure 2 shows the pixelwise confusion of the classes for the trials tested on the multi-class testset. For trials trained on the multi-class trainset IL shows lower confusion with the background for the pancreas and the small intestine than EN. This means pixels of both classes are less often missed to detect. Compared to FS, IL less often confuses the small intestine with the colon. The liver is more often confused with the abdominal wall by IL compared to FS and EN. For the binary trials, IL lowers the confusion of the stomach with

Table 1 Dice score of each class calculated on the testsets of the binary subsets, μ being the mean over all classes

Trainset	Trial	Abd	Col	Liv	Pan	Sma	Sto	μ
Binary	IL (ours)	0.892	0.782	0.741	0.364	0.849	0.709	0.723
	EN (lower)	0.895	0.729	0.745	0.373	0.826	0.503	0.679
MultiClass	FS (upper)	0.764	0.451	0.521	0.170	0.122	0.719	0.458
	IL (ours)	0.730	0.381	0.520	0.184	0.338	0.578	0.445
	EN (lower)	0.673	0.320	0.490	0.139	0.177	0.699	0.416

Every row represents one of the trials, fully supervised (FS), ensemble (EN), and implicit labelling (IL), trained on the binary and multi-class sets. The highest value per column and trainset is highlighted in bold. The classes Abdominal Wall, Colon, Liver, Pancreas, Small Intestine, and Stomach are abbreviated by their first three letters. Cells are highlighted in green if the difference between IL and EN is strongly significant ($p < .01$) or yellow if significant ($p < .05$)

Table 2 Dice score of each class calculated on the testset of the multi-class subset, μ being the mean over all classes

Trainset	Trial	Bac	Abd	Col	Liv	Pan	Sma	Sto	μ
Binary	IL (ours)	0.932	0.578	0.472	0.222	0.503	0.386	0.709	0.543
	EN (lower)	0.955	0.590	0.363	0.375	0.335	0.214	0.503	0.476
MultiClass	FS (upper)	0.970	0.800	0.450	0.862	0.679	0.913	0.719	0.783
	IL (ours)	0.905	0.632	0.251	0.650	0.306	0.455	0.578	0.540
	EN (lower)	0.931	0.728	0.360	0.387	0.239	0.304	0.699	0.512

Every row represents one of the trials, fully supervised (FS), ensemble (EN), and implicit labelling (IL), trained on the binary and multi-class sets. The highest value per column and trainset is highlighted in bold. The classes Background, Abdominal Wall, Colon, Liver, Pancreas, Small Intestine, and Stomach are abbreviated by their first three letters. Cells are highlighted in green if the difference between IL and EN is strongly significant ($p < .01$) or yellow if significant ($p < .05$)

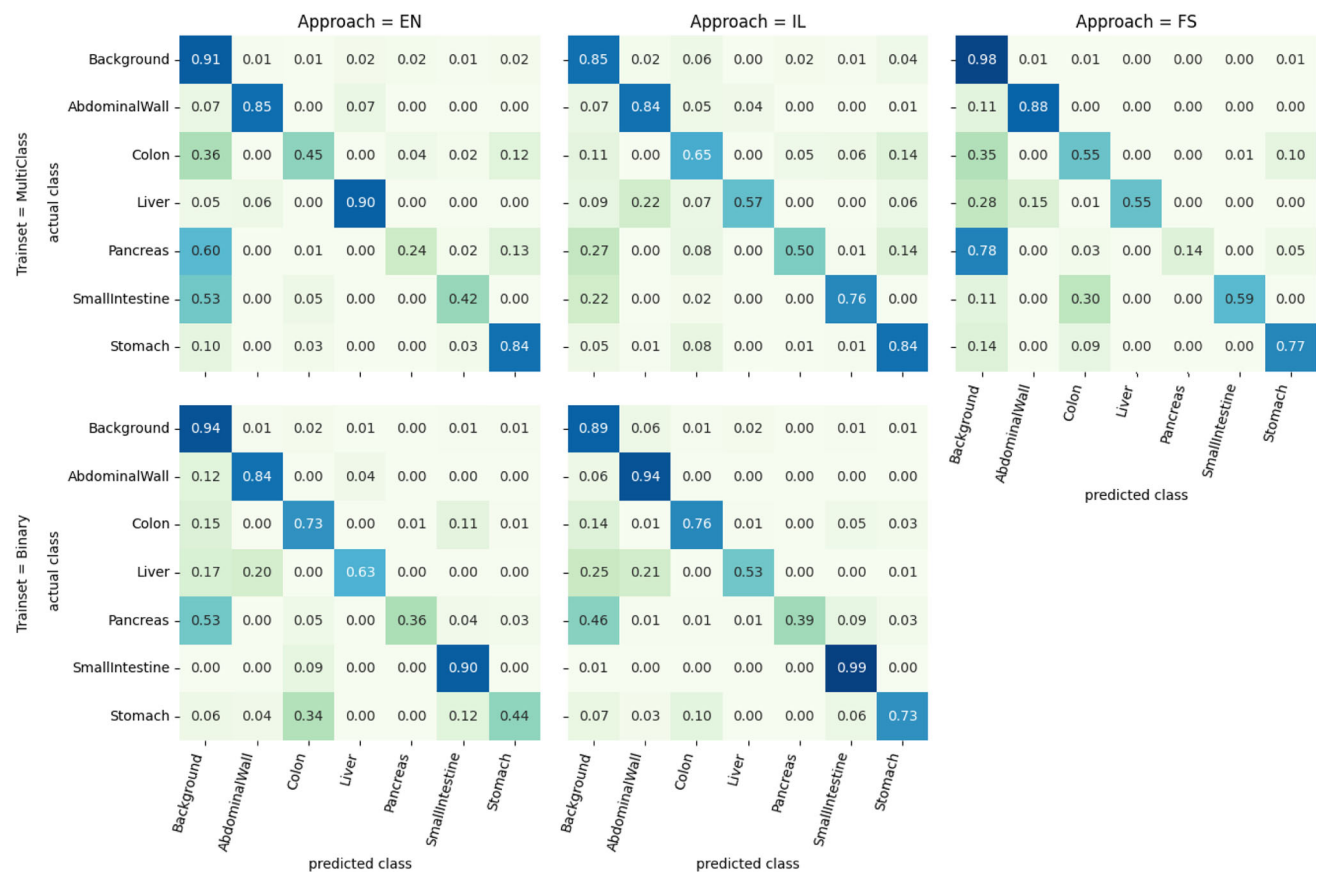


Fig. 2 Pixelwise confusion matrix per approach and trainset tested on the multi-class testset

Fig. 3 Segmentation results for 4 example images. The columns show the original frame, the ground truth, and the predicted segmentations of the fully supervised (FS), implicit labelling (IL), and ensemble (EN) approaches. Rows (a) and (b) show examples of good IL performance; (c) and (d) show more difficult cases

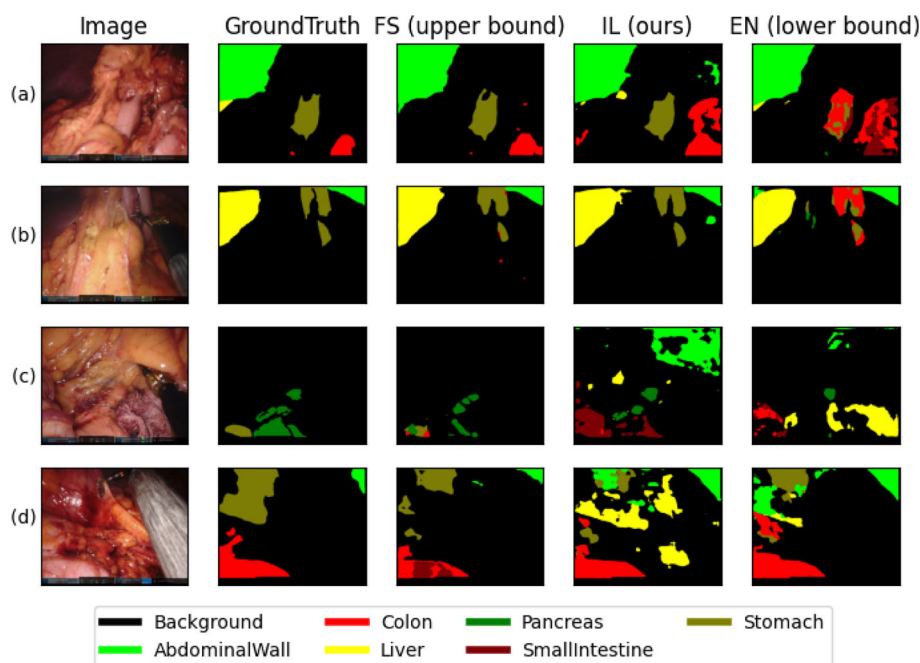


Table 3 Dice score of each class calculated on the testsets of the binary subsets, μ being the mean over all classes

Trial	Abd	Col	Liv	Pan	Sma	Sto	μ
Separate models of EN	0.894	0.790	0.786	0.400	0.866	0.701	0.740
IL (ours)	0.892	0.782	0.741	0.364	0.849	0.709	0.723
IL w/o additional neg. labels	0.897	0.766	0.783	0.326	0.832	0.633	0.706

The first row shows the performance of each separate model in the ensemble before applying the argmax. The last row shows a variant of the IL approach if no additional negative samples are inferred. The classes Abdominal Wall, Colon, Liver, Pancreas, Small Intestine, and Stomach are abbreviated by their first three letters

the colon and small intestine. The confusion of the liver with the abdominal wall and background slightly increases.

Figure 3 shows four examples of segmentation results. In the rows (a) and (b), FS and IL are able to segment a single structure with one class, while the EN approach is mixing multiple classes. In rows (c) and (d) all three approaches produce patched results per organ and detect classes not present.

Table 3 shows the results of our ablation studies trained and evaluated on the binary subset. The first row shows the dice score of the separate models in the ensemble before applying the argmax. All classes except the abdominal wall reach higher scores compared to the merged ensemble. IL is outperformed on all classes except the stomach. The second row shows the results of our IL approach if no additional negative samples are inferred due to mutual exclusion and only the loss masking is applied. While this model is better than our proposed IL approach for the classes abdominal wall and liver, the remaining classes and the average performance are below IL.

The trials were evaluated with respect to the inference time by inferring 1000 random inputs on an Nvidia RTX A5000. The time needed from loading the image to the GPU

to downloading the prediction back to the CPU was averaged. All models were already prepared on the GPU. The ensemble approach required 136 ± 5.8 ms per frame; the fully supervised and implicit labelling approaches both required 23 ± 0.1 ms per frame, resulting in 7fps and 43fps, respectively, and therefore a 6x speed up. Memory consumption on the GPU stayed below 3.2 GB for the ensemble and 2.3 GB for the fully supervised and implicit labelling approaches.

Discussion

In summary, this work demonstrates that the implicit labelling approach is able to leverage multiple complementary datasets into one model. We find that models benefit from a better scene understanding through more learned classes and reach better performance. Further, we point out limitations due to the lack of data diversity we came across in the different data subsets.

As Table 3 shows, the performance on EN merged by argmax drops compared to the performance of its separate models. This again proves the importance of being able to

train a single model being able to understand and segment all classes. Best results are reached by fully annotated datasets, as shown by FS in Tables 1 and 2. But as these are rarely available, our proposed method is a valid alternative in cases where multiple datasets need to be merged.

Tables 1 and 2 clearly show that IL significantly outperforms EN in all four combinations of test and trainset of the binary and multi-class subsets. On the binary train and binary test combination our proposed approach is either within 1% difference on the dice score or significantly better than the lower baseline. This supports our initial assumption that the models benefit from better scene understanding due to complementary information. As shown in Fig. 3(a) and (b) the knowledge of multiple classes, as in the FS and IL approach, helps the model to select one class. This missing information in the EN approach leads to segmentations being patched together of multiple classes. IL is able to solve this issue without the need for fully segmented annotations.

This is demonstrated further in Fig. 2 by comparing the class confusion of IL to EN. For all classes except the liver the confusion drops. In particular, the confusion of colon, small intestine, and stomach on the binary trainset, where those classes do look similar in many cases, is reduced by the knowledge of all classes. Interestingly the fully supervised model is not able to prevent confusion of the small intestine and the colon, which might be caused by rare occurrence of those classes due to the stomach-centric multi-class subset. Interestingly, the large increase of negative samples in the IL approach compared to the FS approach does not lead to the model to prefer the background class; rather, the opposite is the case. This might be due to the higher weighting factors of positive pixels used in the loss. This proves our assumption that models benefit from better scene understanding through more learned classes.

As shown in Tables 1 and 2, the models perform worse when tested on the subset they are not trained on. Especially on the Pancreas and Small Intestine, the FS approach drops in performance. This is explained by the fact that the multi-class subset is based on the stomach. Therefore, all frames contain this structure, limiting the possible viewing angle and distances of other annotated classes. IL is able to maintain the highest cross-domain performance in both directions. This cross-domain capability can further be seen in the first row of Table 2 where IL is trained on the multi-class trainset and tested on the binary subset without any performance drop on the average dice score compared to the third row which shows the performance of IL tested on the multi-class testset. EN, in rows two and four, is dropping by 5% due to the different appearance in the subsets.

The second row of Table 3 shows the importance of the implicit labels, as the model's performance drops without them. This ablation is comparable to the approach presented by Ulrich et al. [10], which further proves the benefit of our

method, as by simply adding the implicit knowledge we outperform related work.

When looking at the situation given by the dataset we are limited to a small fully annotated dataset and have more binary data on hand. While it is possible to train the FS approach on the small dataset and achieve a good performance when testing on similar data, as shown in Table 2, we do see the performance drop when the model is applied to different data, as shown in Table 1. The classes Pancreas and Small Intestine drop by 51% and 79%, respectively, resulting in dice scores below 20%. Considering that the binary dataset is focused around the organ it is tested for, these misclassifications can be considered critical, as the central organ is not recognized. When training our IL approach, on the other hand, not only more data are available, but also the performance drop in changing data is less severe. While we do see dropping performance on classes, if they are no longer in the focus of the frame, no class drops below a dice score of 25%, as shown in Table 2. This proves our IL approach is more applicable in a realistic setting with limited data availability, producing more robust results without the need for complicated preprocessing or much higher computational costs.

Finally, as the inference time results show, the implicit labelling approach combines fast inference with the ability to learn on not fully annotated datasets. The ensemble is 6 times slower, as it needs to infer 6 models instead of one, setting a limit to the scalability. The runtime of the ensemble increases linearly with the number of classes, while our implicit labelling approach allows adding more segmentation data with a negligible effect on runtime.

Conclusion

In this paper, we presented, to the best of our knowledge, the first approach for laparoscopic organ segmentation that combines multiple datasets into one model. We accomplished this by applying a combination of masking during loss and mutual exclusion constraints. We were able to show that segmentation models benefit from a better understanding of the scene in the sense of knowing more classes, improving the overall dice score, reducing confusion between classes, and improving generalization to changes in the appearance of the classes. Further, we were able to show that we do not require all classes to be annotated in a single dataset but rather can combine complementary ones. The resulting model was able to achieve real-time capable inference speeds of 43fps on an Nvidia RTX A5000 GPU, and additional classes can be added without linearly increasing the runtime.

Even though the already good results, we see potential to further improve the approach. For example, the model might include weak labels during training in the form of binary presence of classes. Also, the effects of the use of datasets

with overlapping classes should be investigated, as well as the possibility to combine non-complementary datasets, like different levels of detail in the classes. For example, some datasets only segment instruments in one class, and others split them up into different types. Further, the use of active learning to selectively annotate missing data might be promising, as it matches very well with the structure of complementary datasets. Overall, we see great potential in applying our method in multiple settings of semantic segmentation of surgical data science.

Funding Open Access funding enabled and organized by Projekt DEAL. This work is funded by the German Federal Ministry of Health (BMG), on the basis of a decision by the German Bundestag, within the “Surgomics” project (Grant Number BMG 2520DAT82), the German Cancer Research Center (CoBot 2.0), the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) as part of Germany’s Excellence Strategy (EXC 2050/1, Project ID 390696704) within the Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of the Dresden University of Technology and by the European Union through NEARDATA under the grant agreement ID 101092644. FRK is supported by the Joachim Herz Foundation (Add-On Fellowship for Interdisciplinary Life Science).

Data availability The used DSAD dataset is available at <https://doi.org/10.6084/m9.figshare.21702600>.

Code availability The code is available on https://gitlab.com/nct_tso_public/dsad-segmentation/.

Declarations

Conflict of interest The authors declare no conflict of interest.

Ethical approval For this type of study, formal consent is not required.

Informed consent This article contains patient data from publicly available datasets.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Jin Y, Yu Y, Chen C, Zhao Z, Heng P-A, Stoyanov D (2022) Exploring intra- and inter-video relation for surgical semantic scene segmentation. *IEEE Trans Med Imaging* 41(11):2991–3002. <https://doi.org/10.1109/TMI.2022.3177077>
- Mohammed A, Yildirim S, Farup I, Pedersen M, Hovde Ø (2019) StreoScenNet: surgical stereo robotic scene segmentation. In: *Medical imaging 2019: image-guided procedures, robotic interventions, and modeling*, vol 10951, p 109510. SPIE. <https://doi.org/10.1117/12.2512518>. International Society for Optics and Photonics
- Yoon J, Hong S, Hong S, Lee J, Shin S, Park B, Sung N, Yu H, Kim S, Park S, Hyung WJ, Choi M-K (2022) Surgical scene segmentation using semantic image synthesis with a virtual surgery environment. In: *Medical image computing and computer assisted intervention—MICCAI 2022*. Springer, Cham, pp 551–561
- Fuentes-Hurtado F, Kadhodamohammadi A, Flouty E, Barbarisi S, Luengo I, Stoyanov D (2019) Easylab: weak labels for scene segmentation in laparoscopic videos. *Int J Comput Assist Radiol Surg* 14(7):1247–1257
- Allan M, Kondo S, Bodenstedt S, Leger S, Kadhodamohammadi R, Luengo I, Fuentes F, Flouty E, Mohammed A, Pedersen M, Kori A, Alex V, Krishnamurthi G, Rauber D, Mendel R, Palm C, Bano S, Saibro G, Shih C-S, Chiang H-A, Zhuang J, Yang J, Igloukov V, Dobrenkii A, Reddiboina M, Reddy A, Liu X, Gao C, Unberath M, Kim M, Kim C, Kim H, Lee G, Ullah I, Luna M, Park SH, Azizian M, Stoyanov D, Maier-Hein L, Speidel S (2020) 2018 robotic scene segmentation challenge. <https://doi.org/10.48550/ARXIV.2001.11190>
- HeiChole Surgical Workflow Analysis and Full Scene Segmentation (HeiSurF), EndoVis Subchallenge 2021. <https://www.synapse.org/#!/Synapse:syn25101790/wiki/608802>. Accessed 14 Nov 2022
- Maier-Hein L, Eisenmann M, Sarikaya D, März K, Collins T, Malpani A, Fallert J, Feussner H, Giannarou S, Mascagni P, Nakawala H, Park A, Pugh C, Stoyanov D, Vedula SS, Cleary K, Fichtinger G, Forestier G, Gibaud B, Grantcharov T, Hashizume M, Heckmann-Nötzel D, Kennigott HG, Kikinis R, Mündermann L, Navab N, Onogur S, Roß T, Sznitman R, Taylor RH, Tizabi MD, Wagner M, Hager GD, Neumuth T, Padoy N, Collins J, Gockel I, Goedeke J, Hashimoto DA, Joyeux L, Lam K, Leff DR, Madani A, Marcus HJ, Meireles O, Seitel A, Teber D, Ückert F, Müller-Stich BP, Jannin P, Speidel S (2022) Surgical data science: from concepts toward clinical translation. *Med Image Anal* 76:102306. <https://doi.org/10.1016/j.media.2021.102306>
- Carstens M, Rinner FM, Bodenstedt S, Jenke AC, Weitz J, Distler M, Speidel S, Kolbinger FR (2023) The Dresden surgical anatomy dataset for abdominal organ segmentation in surgical data science. *Sci Data* 10(1):1–8. <https://doi.org/10.1038/s41597-022-01719-2>
- Shi G, Xiao L, Chen Y, Zhou SK (2021) Marginal loss and exclusion loss for partially supervised multi-organ segmentation. *Med Image Anal* 70:101979. <https://doi.org/10.1016/j.media.2021.101979>
- Ulrich C, Isensee F, Wald T, Zenk M, Baumgartner M, Maier-Hein KH (2023) Multitalent: a multi-dataset approach to medical image segmentation. In: *Medical image computing and computer assisted intervention: MICCAI 2023*. Springer, Cham, pp 648–658
- Dmitriev K, Kaufman AE (2019) Learning multi-class segmentations from single-class datasets. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp 9501–9511
- Yan K, Cai J, Zheng Y, Harrison AP, Jin D, Tang Y, Tang Y, Huang L, Xiao J, Lu L (2020) Learning from multiple datasets with heterogeneous and partial labels for universal lesion detection in CT. *IEEE Trans Med Imaging* 40(10):2759–2770
- Dice LR (1945) Measures of the amount of ecologic association between species. *Ecology* 26(3):297–302. <https://doi.org/10.2307/1932409>
- Kolbinger FR, Rinner FM, Jenke AC, Carstens M, Krell S, Leger S, Distler M, Weitz J, Speidel S, Bodenstedt S (2023) Anatomy segmentation in laparoscopic surgery: comparison of machine learning and human expertise—an experimental study. *Int J Surg* 109(10):2962–2974. <https://doi.org/10.1097/JS9.0000000000000595>

15. Chen L-C, Papandreou G, Schroff F, Adam H (2017) Rethinking atrous convolution for semantic image segmentation. <https://doi.org/10.48550/arXiv.1706.05587>
16. Lin T-Y, Maire M, Belongie S, Bourdev L, Girshick R, Hays J, Perona P, Ramanan D, Zitnick CL, Dollár P (2014) Microsoft COCO: common objects in context. <https://doi.org/10.48550/arxiv.1405.0312>
17. Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Köpf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintala S (2019) PyTorch: an imperative style, high-performance deep learning library. <https://doi.org/10.48550/arxiv.1912.01703>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.