

Addressing multi-label imbalance problem of surgical tool detection using CNN

Manish Sahu¹  · Anirban Mukhopadhyay¹ · Angelika Szengel¹ · Stefan Zachow¹

Received: 1 February 2017 / Accepted: 14 March 2017 / Published online: 29 March 2017
© CARS 2017

Abstract

Purpose A fully automated surgical tool detection framework is proposed for endoscopic video streams. State-of-the-art surgical tool detection methods rely on supervised *one-vs-all* or *multi-class* classification techniques, completely ignoring the co-occurrence relationship of the tools and the associated class imbalance.

Methods In this paper, we formulate tool detection as a *multi-label* classification task where tool co-occurrences are treated as separate classes. In addition, imbalance on tool co-occurrences is analyzed and stratification techniques are employed to address the imbalance during convolutional neural network (CNN) training. Moreover, temporal smoothing is introduced as an online post-processing step to enhance runtime prediction.

Results Quantitative analysis is performed on the *M2CAI16 tool detection dataset* to highlight the importance of stratification, temporal smoothing and the overall framework for tool detection.

Conclusion The analysis on tool imbalance, backed by the empirical results, indicates the need and superiority of the proposed framework over state-of-the-art techniques.

Keywords Transfer learning · Surgical tool detection · CNN · Laparoscopic videos · Multi-label learning

Introduction

Fast and accurate recognition of surgical workflow plays an important role in modern computer-assisted intervention

(CAI). Modern operating rooms (OR) demand monitoring of surgical processes to reduce preventable errors, the absence of which may result in failures up to the loss of human lives [5]. In addition, multitudes of other OR procedures, for example automated clinical assistance, staff assignment, etc., can benefit from surgical workflow recognition [1, 2, 17]. Recent CAI literature [3, 9, 17, 19] has identified that surgical tool occurrences are closely related to the phases of surgical workflow. Moreover, tool detection and tracking on endoscopic images has the potential of controlling a robot-mounted endoscopic camera holder, especially for solo surgeries [11, 18]. Change in illumination, specular reflection and partial occlusion are some of the major challenges that render surgical tool detection a challenging task. This work mainly focuses on a fully automatic identification of surgical tool(s) from endoscopic video streams.

State-of-the-art methods treat surgical tool detection as a supervised *multi-class* or *one-vs-all* classification task [17]. Based on the observation that multiple tool co-occurrences happen quite often in endoscopic video frames, we have formulated the task as a generalized *multi-label* classification. The co-occurrence of tools, formally termed as *label-sets*, forms the output set of *multi-label* classification. In particular, rather than treating each tool as a stand-alone class, we've considered *label-sets* of multiple tools, along with the introduction of a no-tool (i.e., background) class. Second-order co-occurrences of tools along with frequency of occurrences are visualized in Fig. 1 for intuitive understanding. Interesting tool co-occurrence relationship patterns emerge from such visualization. For example, though *hook* is the most often used tool in surgical intervention, in second order, it is used only in association with grasper. Other interesting relationships involving two-tool usage can also be inferred from Fig. 1.

✉ Manish Sahu
sahu@zib.de

¹ Zuse Institute Berlin, Berlin, Germany

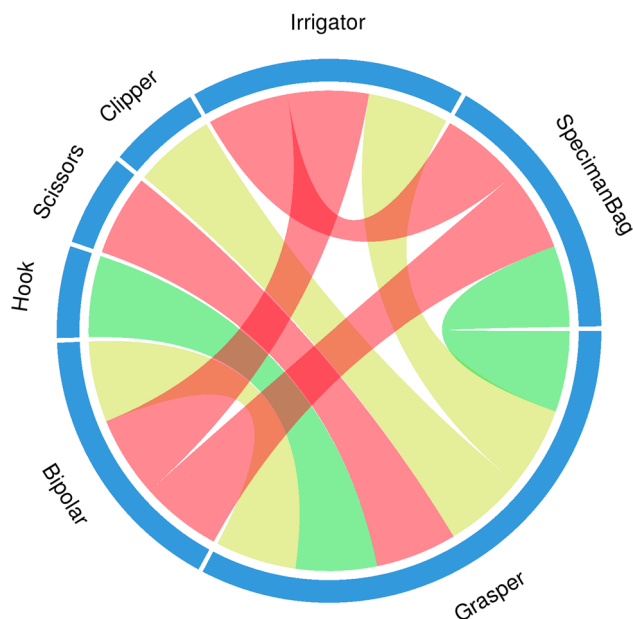


Fig. 1 Chord diagram [6] showing the second-order co-occurrences of tools (two tools together) in *M2CAI16 tool detection training dataset*. Tool usage frequency is color-coded in a discrete fashion: *red* (0–160), *yellow* (160–1000) and *green* (>1000)

One key observation of this work is the imbalance associated with the tool usage (could be understood intuitively from Fig. 1). It has already been proven that the imbalance in data affects the *binary* and *multi-class* classification accuracy [4]. However, the imbalance associated with the tool usage during a surgery has not been quantified before. In this work, imbalance in tool usage during an endoscopic intervention is analyzed quantitatively using the measures introduced by Chartre et al. [4]. Moreover, the effects of tool usage imbalance on detection accuracy are quantitatively analyzed in a novel experimental setup and specific sampling strategy to address imbalance [13] in tool usage is introduced during CNN training.

Major contributions of this work are twofold. First, surgical tool detection is performed in a generalized *multi-label* setting, with quantitative analysis focusing on handling surgical tool imbalance. To the best of our knowledge, this is the first work where tool detection is formulated as a general *multi-label* classification problem. Secondly, a novel transfer learning architecture is proposed for fine tuning and domain adaptation of AlexNet [7] toward a surgical tool identification task. In particular, weighted *uni-variate* loss (as learning objective) for joint output distribution is adopted for handling residual tool imbalance after stratification.

Related work

Surgical tool detection within endoscopic videos has received increasing attention in recent years [3, 15]. For the sake of

brevity, we have mainly considered works where endoscopic video streams are considered as the sole input source. Tool detection procedures from video streams often consider the task as both tool identification and localization problem [3]. For example, Speidel et al. [14] suggested a tool identification pipeline that consists of segmentation and 3D model-based processing. Sahu et al. [11] proposed detection and tracking of surgical tools over the virtual control interface on endoscopic stream to control a robot.

This work, however, similar to Twinanda et al. [17], considers tool detection as a tool presence detection task without explicit localization. For the rest of the paper, “tool detection” is commonly used to refer to automatic tool identification from endoscopic video frames. Twinanda et al. [17] proposed deep learning-based features to be used in *one-vs-all* classification framework for tool identification without localization. Most recently, performance of different tool detection techniques [10, 12, 16] is quantitatively evaluated in *M2CAI16 tool detection challenge*. In particular, Twinanda et al. [16] used ToolNet—a network very similar to AlexNet with a final tool detection layer. Raju et al. [10] used an ensemble of two networks and Sahu et al. [12] consisted of a *multi-label* learning approach with no stratification, followed by random forest for classification. In this work, unlike other proposed techniques, we have generalized the problem as *multi-label* classification task, quantitatively analyzed the imbalance and adopted strategies to overcome imbalance-related issues.

Method

In this section, we first provide an overview of *multi-label* classification for tool detection task. Next, we describe metrics used for defining levels of imbalance in a *multi-label* dataset in “Imbalance quantification” section, followed by a description of stratification technique used to address *multi-label* imbalance in “Stratification” section. In “ZIBNet” section, ZIBNet architecture is introduced with novel design choices that are incorporated during learning. Finally, we propose temporal smoothing as an online post-processing step which suppresses false positives during prediction.

Multi-label classification

Supervised classification-based surgical tool detection have focused on formulating the problem in a *one-vs-all* or *multi-class* setting. Even though simple and intuitive, in this paper we argue that these settings do not address the problem in its general sense. In particular, due to the co-occurrence of multiple surgical tools at different endoscopic video frames, general *multi-label* classification should be used to model tool presence instead. *Multi-label* classification is the generalization of *binary* or *multi-class* classification. In this




Fig. 2 *Upset* [8] visualization of the third-order co-occurrence (three tools appearing together) of the tools in *M2CAI16 tool detection training* dataset. It shows possible three tool combinations (*bottom right*) with orange and black color representing presence and absence of these

combinations in the training dataset respectively. The third-order co-occurrences of the tools are shown on top with corresponding individual tool distribution on *bottom left*

scenario, no a priori limit on the number of tools present in the output set is imposed during classification.

For example, second-order co-occurrences of surgical tools (i.e., two tools appearing together) are reported in Fig. 1. In *one-vs-all* or *multi-class* classification setting, the desired chord diagram would be an empty circle, penalizing all the interconnected entries. However, plotting the ground truth annotation in Fig. 1 suggests the existence of interconnected entries (more than one tool), which would have been penalized in the earlier settings. Similarly, third-order tool co-occurrence with respect to instrument distribution is visualized in Fig. 2 which shows distribution of three tools appearing together. The presence of the co-occurrences in Figs. 1 and 2 violates the mutually exclusive class assumption of *one-vs-all* or *multi-class* classification. This motivates us to model the problem as a *multi-label* classification one where co-occurrence entries are also considered valid and not penalized during classification.

Formally, for all N annotated video frames in our training dataset $F \in \{f_i\}$ where $i = 1, 2, \dots, N$, a *multi-label* classifier C learns to represent the total set of tool labels $T \in \{t_j\}$ where $j = 1, 2, \dots, M$. For a testing video stream, C must produce as output a set $Z_i \subseteq T$ with predicted tool labels for the i -th video frame. Note that, this generalization results in 2^M potential combinations, which are termed as label-sets. In a general setting, this might result in many prac-

tical constraints (e.g., memory for storing, representation and performing actual classification). However due to the practicalities of endoscopic intervention, where only a limited number of tools (maximum three for this particular dataset) and combinations (see Figs. 1 and 2) can be present at once, the *multi-label* problem remains tractable. In particular, the seven-tool *M2CAI16 tool detection* dataset has resulted in approximately twenty *label-sets* spanning from order zero to three.

Imbalance quantification

Even though imbalance in *binary* and *multi-class* classifications is a well-studied problem, quantitative analysis of imbalance in *multi-label* dataset is proposed in very few occasions [4, 13]. Conventional imbalance analysis methods, designed for *binary/multi-class* classification, assume only the ratio of majority to minority class labels as imbalance measure, therefore, not suitable for *multi-label* datasets. There exist some traditional metrics notably label cardinality and label density which characterize *multi-label* datasets. Label cardinality is the average number of active labels per sample, and label density denotes average of label cardinality over the total number of labels. Mathematically, these measures can be defined as follows [4]:

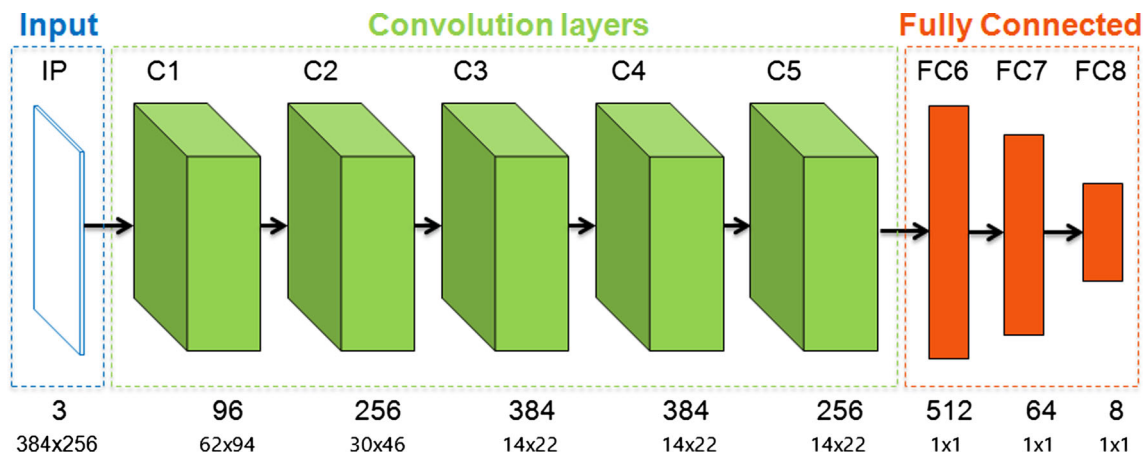


Fig. 3 Proposed CNN architecture. The input layer (blue) size and fully connected layers (orange) are adapted to tool detection task, while the convolution layers are similar to the AlexNet [7] architecture

$$\text{Cardinality}(F) = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M t_{ij} \quad (1)$$

$$\text{Density}(F) = \frac{\text{Cardinality}(F)}{M} \quad (2)$$

With the presence of multiple *label-sets*, special metrics are needed for analyzing imbalance in *multi-label* datasets in detail. In this work, we have exploited imbalance ratio per label (IRLbl), the mean imbalance ratio (MeanIR) and the coefficient of variation of IRLbl (CVIR) introduced in [4] to analyze the imbalance in our tool dataset:

$$\text{MeanIR} = \frac{1}{M} \sum_{i=1}^M \text{IRLbl}(t_i) \quad (3)$$

$$\text{CVIR} = \text{MeanIR} \sqrt{\frac{\sum_{i=1}^M (\text{IRLbl}(t_i) - \text{MeanIR})^2}{M - 1}} \quad (4)$$

Here IR per label, $\text{IRLbl}(T_i)$, is calculated as the ratio between the majority (most frequent) label and label T_i . As a result, the majority label will always have $\text{IRLbl} = 1$ and rest of the labels will have higher $\text{IRLbl} > 1$. MeanIR computes the average level of imbalance of the dataset, while CVIR measures variation of IRLbl, i.e., similarity of level of imbalance between all labels. For a perfectly balanced dataset, all IRLbl values would be 1, which results in values of MeanIR and CVIR being 1 and 0, respectively. The joint use of MeanIR and CVIR with values greater than 1 and 0, respectively, denotes the level of imbalance in a *multi-label* dataset. Moreover, the values of IRLbl greater than 1 can be used for measuring individual label imbalance.

Stratification

Stratification is the process of sampling, where the proportion of disjoint groups is maintained [13]. Stratification in the *multi-label* data context is a challenging task. Improper stratification might significantly reduce the performance of classifiers as demonstrated by Sechidis et al. [13].

The most intuitive stratification in the tool detection setting would be to consider a *balanced* strategy. In this setting, the occurrence frequency of the least frequent tool would be considered as the desired sample size and the rest of the tools would be sampled accordingly. However, co-occurrence of tools in different frames actually results in an unbalanced training and validation set.

A better way of handling the problem would be to consider stratification on *label-sets*. Here, the frequency of *label-sets* is considered for sampling of image frames. A stratification threshold ψ is applied, where for more frequent *label-sets*, ψ occurrences are sampled randomly, and for less frequent ($< \psi$) *label-sets*, all samples are considered for training.

ZIBNet

Our proposed CNN architecture (Fig. 3) is composed of three main parts:

- *Input layer* which accepts an input image of size 384×256 pixels.
- *Convolutional layers* which are similar to AlexNet architecture.
- *Fully connected layers* which are specific to tool detection task with size 512, 64 and 8, respectively.

Before the learning step, the convolutional layers weights are initialized with AlexNet convolutional weights and the

fully connected layers are initialized with random weights. The rectify units are applied to the output of convolutional and fully connected layers except the last fully connected layer which is connected to sigmoid nonlinearity. Since our network contains pre-trained convolutional weights of AlexNet which are generic for earlier convolution layers and become specific to ImageNet objects for higher layers, we assign layer-specific learning rates

$$\text{LR} = c_i \eta \quad (5)$$

where η is the learning rate and $c_i = [0.1, 0.2, 0.3, 0.4, 0.5, 1.0, 1.0, 1.0]$ is the learning coefficient for each layer. The learning coefficient becomes higher with subsequent convolutional layer, except for the fully connected layers whose coefficients remain fixed since these layers have been randomly initialized. The output layer contains eight units for seven tools and an additional “no-tool” label, which is added to the ground truth and represents that none of the given tools are present in the image, i.e., background.

During the learning step, the network minimizes the joint label distribution through a *uni-variate* loss function \mathcal{L} defined as:

$$\mathcal{L} = - \sum_{t=1}^T w(t) [z_t \log y_t + (1 - z_t) \log(1 - y_t)] \quad (6)$$

where $w(t)$ is a weighting function that normalizes loss in terms of output z_t , and y_t is prediction for tool t . Intuitively, even after *label-set* stratification, imbalance on the tool label would not be omitted completely. \mathcal{L} is formulated as a *uni-variate* loss function where cross entropy is weighted with tool occurrence frequency in the training data, to manage the residual imbalance after stratification.

To avoid over-fitting, we perform real-time data augmentation (flipping, mirroring and cropping) during learning, apply dropout units for “FC6” and “FC7” layers and use a weight decay of (10^{-5}) for every layer. Finally, the network is trained using stochastic gradient descent (η of (10^{-2})) with momentum of 0.9 using the holdout scheme.

Temporal smoothing

Due to the stochastic nature of the classification process, false detections would occur during the testing step. For reducing such false detections, we have adopted a temporal smoothing (TS) approach as an online post-processing step. It assumes that each tool transition within the endoscopic videos is smooth and takes previous frame detections into account in a weighted scheme. A window of five frames (including current and four previous frames) with normalized linear weights determines the current output detection. Mathematically, TS is defined as follows:

$$y_{ts} = \frac{\sum_{i=0}^t w_i y_{t-i}}{\sum_{i=0}^t w_i} \quad (7)$$

where y_{ts} is the temporally smooth output for the current time step, $t = 4$ and $w_i = [1.0, 0.8, 0.6, 0.4, 0.2]$ is the weight coefficient for each time step.

Results

This section provides a quantitative analysis of the proposed method, as well as quantitative comparison w.r.t. state-of-the-art methods, to demonstrate its effectiveness for surgical tool detection from endoscopic video streams.

Data preparation

All our quantitative experiments were performed on the *M2CAI Tool Detection Challenge 2016* training and testing dataset [16]. The dataset contains 15 clinical cholecystectomy surgical procedures (one video per surgical procedure), which were performed using seven surgical tools (as shown in Fig. 4). The dataset is divided into 10 training and 5 testing videos. The videos were captured at 25 frames per second, but the ground truth (GT) annotation was done at 1 frames per second. A tool was considered to be present only if at least half of the tool tip was visible.

To provide a fair comparison, we only considered average precision (AP) per tool and mean average precision (mAP) for all tools [17] as the comparison metric. For all




Fig. 4 Appearances of different surgical tools from *M2CAI16 tool detection dataset*. The tools used during the surgical procedure (left to right) are grasper, bipolar, hook, clipper, scissors, irrigator and specimen bag

Table 1 Comparison of meanAP for state-of-the-art techniques on *M2CAI tool detection testing* dataset

	ToolNet [16]	Sahu et al. [12]	Raju et al. [10]	Proposed
MeanAP	0.52	0.53	0.63	0.65

the experiments, *M2CAI tool detection training* dataset was used for training and the results were reported on *M2CAI Tool Detection Testing* dataset with stratification threshold $\psi = 300$.

Comparison with state-of-the-art methods

The proposed framework (*label-set* based stratification and weighted *uni-variate* loss for training ZIBNet followed by TS) results in state-of-the-art performance when relying on mAP for evaluation. We compared the results of our proposed method with the top performing methods from *M2CAI Tool Detection Challenge*, all of which [10, 12, 16] relied on CNN for tool detection. All the results are taken from *M2CAI Tool Detection Challenge* website where the respective authors used their own software for training and testing. As Table 1 shows, our proposed method outperformed all other methods. In particular, ZIBNet significantly outperformed [16, 17] in a paired *t* test with $p < 0.05$. We attribute the superior performance of our proposed method to our particular formulation of the problem as a *multi-label* one and our novel stratification technique.

For a detailed performance evaluation of our proposed framework, a receiver operating characteristic (ROC) curve for each tool is shown in Fig. 5. Apart from *Irrigator* and *Scissors*, all the other tools performed sufficiently well. Difficulties in detecting *Scissors* is already known [16, 17]; however, the reason behind low performance on *Irrigator*

Table 2 Imbalance ratio per label (IRLbl) of various tools in the training and testing dataset

Tools	Training	Testing
Grasper	1.28	1.13
Bipolar	22.25	20.76
Hook	1.00	1.00
Scissors	34.37	47.17
Clipper	16.09	23.66
Irrigator	14.82	46.88
Specimen Bag	9.39	15.43

is analyzed in the remainder of this work (“Discussion and conclusion” section).

Imbalance analysis

We have performed a detailed imbalance analysis of all the tools present in the training and testing dataset. Rather than reporting the exact number of occurrences for each tool, we concentrated on the IRLbl measure for each tool (Table 2). As described in “Imbalance quantification” section, the value of the most frequent tool is 1 and rest have higher values (ideally 1, higher IRLbl means higher imbalance for the associated tool). As we can see in Table 2, the most frequent tool is *Hook* with IRLbl 1.0. It is interesting to note that *Irrigator* is almost *three times* more frequent in training dataset compared to Testing dataset. The effects of imbalance are further reported through other metrics in Tables 2 and 3.

Traditionally, for *one-vs-all* or *multi-class* classification, imbalance is reported over the whole dataset using cardinality and density, as shown in Eqs. (1) and (2). However, it is evident from the Eqs. (3) and (4) that overall imbalance of a *multi-label* dataset can be appreciated only by looking at


Fig. 5 Receiver operating characteristic (ROC) curve for presence detection of each tool

Table 3 Values of traditional metrics (label cardinality and label density) and specific metrics (MeanIR and CVIR) over the whole training and testing dataset

Measures	Training	Testing
Cardinality	1.26	1.23
Density	0.18	0.17
MeanIR	14.17	22.29
CVIR	0.83	0.85

Table 4 Quantitative comparison between different imbalance handling (stratification) strategies: no sampling (Unbalanced), tool-level sampling (tool-balanced) and label-set-based sampling (label-set) for various tools

Tools	Unbalanced	Tool-balanced	Label-set
Grasper	0.87	0.61	0.81
Bipolar	0.21	0.46	0.57
Hook	0.96	0.95	0.96
Scissors	0.09	0.3	0.34
Clipper	0.34	0.51	0.46
Irrigator	0.21	0.21	0.24
Specimen bag	0.81	0.64	0.71
MeanAP	0.50	0.53	0.59

MeanIR and CVIR together. Imbalance in the training dataset resulted in MeanIR value of approximately 14 times higher than ideal (ideally 1) with 83% variance (ideally 0%) in IRLbl values as reported in Table 3.

Analysis of stratification techniques

We performed quantitative comparison of different sampling approaches to highlight the importance of stratification on the performance of ZIBNet for tool detection. *Label-sets*-based stratification (“Stratification” section) is compared to unbalanced sampling and tool-level balanced stratification (“Stratification” section). Average precision for each tool and overall mAP is reported in Table 4. In particular, the baseline strategy (no stratification at all)—termed as “Unbalanced” in Table 4, performed worst. Tool-level balanced stratification resulted in an overall increase of 3% over “Unbalanced,” whereas the proposed *label-set* stratification increased mAP by 9%, as shown in Table 4. It is worth noting that Unbalanced approach favored most frequent tools. Stratification, on the other hand, adapted the network to alleviate this bias and enhanced detection of less frequent tools.

Analysis of temporal smoothing

The last part of our design is to apply TS to decrease stand-alone activations as described in “Temporal smoothing” section. We have run experiments to quantitatively demon-

Table 5 Importance of temporal smoothing (TS) on overall tool detection accuracy for different tools over label-based sampling (balanced with TS) and *label-set*-based sampling (label-set with TS)

Tools	Balanced with TS	Label-set with TS
Grasper	0.62	0.83
Bipolar	0.58	0.66
Hook	0.97	0.97
Scissors	0.47	0.51
Clipper	0.65	0.54
Irrigator	0.21	0.29
Specimen bag	0.73	0.77
MeanAP	0.61	0.65

strate the importance of TS. TS is independent of the rest of the proposed framework and as shown in Table 5, consistently improved the performance of both stratification techniques by reducing false positives. In particular, the simple stratification benefited more from TS with an overall mAP increase of 8%, whereas 6% boost in performance was observed for *label-set* stratification.

Discussion and conclusion

Detection of surgical tools in endoscopic video is an important problem requiring a rigorous understanding of the data as well as an effective handling approach. A successful surgical tool detection technique can potentially improve a multitude of CAI applications. For example, detection of surgical workflow phases can directly benefit from tool detection results. However, tool co-occurrences, change of illumination, specular reflection and partial occlusion make the detection problem significantly more difficult.

This work clearly showed that generalizing the classification problem with domain adaptation can significantly improve classification results. Our proposed method demonstrates that fully automatic tool detection results in an acceptable level of agreement with the manual annotations. By modeling tool co-occurrences as *label-sets*, we can better handle the inherent structure of surgical tool presence during interventions. Moreover, a detailed study of imbalance in *label-sets* has motivated us to develop stratification methods for CNN training.

Note that, IRLbl in Table 2 suggests that *irrigator* has significantly different imbalance ratio between training and testing dataset. Not only our results consistently lead to lowest AP as reported in Tables 4 and 5, ToolNet results by Twinanda et al. [16] also reported the same. This suggests that along appearance difficulties (in case of *scissors*), label imbalance significantly challenges the performance of CNN.

An important observation of our study is the boost in performance by addition of temporal smoothing as an online post-processing method (no future information is considered). TS consistently enhances detection results for both stratification approaches as reported in Table 5.

This study solely concentrated on tool presence identification; however, future studies of surgical workflow phase recognition can benefit from the insights. In particular, the correlation of surgical phases with the tools being used therein can be exploited further in the *label-set* setting. The imbalance of *label-sets* also suggests special tool co-occurrences which could be used as important phase-transition cues.

In conclusion, this study motivates us to rethink about the standard assumptions regarding surgical tool presence detection. Deviating from de facto supervised *one-vs-all* or *multi-class* techniques (the performance of which heavily depends on the co-occurrence frequencies) toward *multi-label* settings can provide multiple benefits. Finally such fully automatic techniques are expected to be instrumental in advancing the computer assistance during surgical intervention.

Acknowledgements This study was funded by German Federal Ministry of Education and Research (BMBF) under the project BIOPASS (Grant No. 16 5V 7257).

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Human and animal rights statement This article does not contain any studies with human participants or animals performed by any of the authors.

Informed consent This article contains patient data from a publically available dataset.

References

- Allan M, Chang PL, Ourselin S, Hawkes DJ, Sridhar A, Kelly J, Stoyanov D (2015) Image based surgical instrument pose estimation with multi-class labelling and optical flow. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 331–338
- Blum T, Feußner H, Navab N (2010) Modeling and segmentation of surgical workflow from laparoscopic video. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 400–407
- Bouget D, Benenson R, Omran M, Riffaud L, Schiele B, Jannin P (2015) Detecting surgical tools by modelling local appearance and global shape. *IEEE Trans Med Imaging* 34(12):2603–2617
- Charte F, Rivera AJ, del Jesus MJ, Herrera F (2015) Addressing imbalance in multilabel classification: measures and random resampling algorithms. *Neurocomputing* 163:3–16
- Donaldson MS, Corrigan JM, Kohn LT (2000) To err is human: building a safer health system, vol 6. National Academies Press, Washington
- Gu Z, Gu L, Eils R, Schlesner M, Brors B (2014) circlize implements and enhances circular visualization in R. *Bioinformatics*. Oxford Univ Press, p btu393
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems, pp 1097–1105
- Lex A, Gehlenborg N, Strobel H, Vuillemot R, Pfister H (2014) Upset: visualization of intersecting sets. *IEEE Trans Visual Comput Graphics* 20(12):1983–1992
- Padoy N, Blum T, Ahmadi SA, Feussner H, Berger MO, Navab N (2012) Statistical modeling and recognition of surgical workflow. *Med Image Anal* 16(3):632–641
- Raju A, Wang S, Huang J (2016) M2cai surgical tool detection challenge report. <http://camma.u-strasbg.fr/m2cai2016/reports/Raju-Tool.pdf>
- Sahu M, Moerman D, Mewes P, Mountney P, Rose G (2016a) Instrument state recognition and tracking for effective control of robotized laparoscopic systems. *Int J Mech Eng Rob Res* 5(1):33
- Sahu M, Mukhopadhyay A, Szengel A, Zachow S (2016b) Tool and phase recognition using contextual CNN features. arXiv preprint [arXiv:1610.08854](https://arxiv.org/abs/1610.08854)
- Sechidis K, Tsoumakas G, Vlahavas I (2011) On the stratification of multi-label data. In: Joint European conference on machine learning and knowledge discovery in databases, Springer, pp 145–158
- Speidel S, Benzko J, Krappe S, Sudra G, Azad P, Müller-Stich BP, Gutt C, Dillmann R (2009) Automatic classification of minimally invasive instruments based on endoscopic image sequences. In: SPIE medical imaging, International society for optics and photonics, p 72,610A
- Sznitman R, Becker C, Fua P (2014) Fast part-based classification for instrument detection in minimally invasive surgery. In: International conference on medical image computing and computer-assisted intervention, Springer, pp 692–699
- Twinanda AP, Mutter D, Marescaux J, de Mathelin M, Padoy N (2016a) Single- and multi-task architectures for tool presence detection challenge at M2CAI 2016. arXiv preprint [arXiv:1610.08851](https://arxiv.org/abs/1610.08851)
- Twinanda AP, Shehata S, Mutter D, Marescaux J, de Mathelin M, Padoy N (2016b) Endonet: a deep architecture for recognition tasks on laparoscopic videos. arXiv preprint [arXiv:1602.03012](https://arxiv.org/abs/1602.03012)
- Voros S, Long JA, Cinquin P (2007) Automatic detection of instruments in laparoscopic images: a first step towards high-level command of robotic endoscopic holders. *Int J Rob Res* 26(11–12):1173–1190
- Zappella L, Béjar B, Hager G, Vidal R (2013) Surgical gesture classification from video and kinematic data. *Med Image Anal* 17(7):732–745