



# **SurgiTrack: Fine-grained multi-class multi-tool tracking in surgical videos**

Chinedu Innocent Nwoye, Nicolas Padoy

## **► To cite this version:**

Chinedu Innocent Nwoye, Nicolas Padoy. SurgiTrack: Fine-grained multi-class multi-tool tracking in surgical videos. *Medical Image Analysis*, 2025, 101, pp.103438. <10.1016/j.media.2024.103438>. <hal-04872462>

**HAL Id: hal-04872462**

**<https://hal.science/hal-04872462v1>**

Submitted on 8 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License



# SurgiTrack: Fine-grained multi-class multi-tool tracking in surgical videos

Chinedu Innocent Nwoye<sup>\*,</sup>, Nicolas Padoy<sup>\*,</sup>

University of Strasbourg, CAMMA, ICube, CNRS, INSERM, France  
IHU Strasbourg, Strasbourg, France

## ARTICLE INFO

### Keywords:

Surgical tool tracking  
Tool detection  
Surgeon operator  
Multi-class multi-object tracking

## ABSTRACT

Accurate tool tracking is essential for the success of computer-assisted intervention. Previous efforts often modeled tool trajectories rigidly, overlooking the dynamic nature of surgical procedures, especially tracking scenarios like out-of-body and out-of-camera views. Addressing this limitation, the new CholecTrack20 dataset provides detailed labels that account for multiple tool trajectories in three *perspectives*: (1) intraoperative, (2) intracorporeal, and (3) visibility, representing the different types of temporal duration of tool tracks. These fine-grained labels enhance tracking flexibility but also increase the task complexity. Re-identifying tools after occlusion or re-insertion into the body remains challenging due to high visual similarity, especially among tools of the same category. This work recognizes the critical role of the tool operators in distinguishing tool track instances, especially those belonging to the same tool category. The operators' information are however not explicitly captured in surgical videos. We therefore propose SurgiTrack, a novel deep learning method that leverages YOLOv7 for precise tool detection and employs an attention mechanism to model the originating direction of the tools, as a proxy to their operators, for tool re-identification. To handle diverse tool trajectory perspectives, SurgiTrack employs a harmonizing bipartite matching graph, minimizing conflicts and ensuring accurate tool identity association. Experimental results on CholecTrack20 demonstrate SurgiTrack's effectiveness, outperforming baselines and state-of-the-art methods with real-time inference capability. This work sets a new standard in surgical tool tracking, providing dynamic trajectories for more adaptable and precise assistance in minimally invasive surgeries.

## 1. Introduction

Surgical tool tracking plays a pivotal role in computer-assisted surgical systems, offering valuable insights for a range of applications, including skill assessment (Pedrett et al., 2023), visual servoing (Xu et al., 2023), navigation (Xu et al., 2022), laparoscope positioning (Dutkiewicz et al., 2005), safety and risk zone estimation (Richa et al., 2011), and augmented reality (Martin-Gomez et al., 2023). While tool detection identifies target tools in an image or frame, tool tracking goes a step further to also include the estimation and prediction of the tools' locations as they appear in subsequent video frames. Historically, tool tracking relied on traditional machine learning features, encompassing color, texture, SIFT, and geometry (Pezementi et al., 2009; Sznitman et al., 2012; Alsheakhali et al., 2015; Dockter et al., 2014; Du et al., 2016). Recent advances in deep learning (Bouget et al., 2017; Lee et al., 2019; Nwoye et al., 2019; Zhao et al., 2019b,a; Robu et al., 2021; Nwoye, 2021; Fathollahi et al., 2022; Wang et al., 2022; Rueckert et al., 2023) ushered in a new era by enabling the extraction of more robust features for tool re-identification (re-ID). Despite the remarkable progress, challenges remains. Existing efforts have primarily focused on

single tool tracking (Zhao et al., 2019b), single-class multi-tool tracking (Fathollahi et al., 2022), or multi-class single-tool tracking (Nwoye et al., 2019). Whereas in real-world surgical scenarios, multiple tools of varying classes are often utilized simultaneously, necessitating multi-class multi-tool tracking, a domain that remains largely unexplored mainly due to lack of the requisite dataset.

Recently, a new dataset known as CholecTrack20 (Nwoye et al., 2023) was introduced providing the multi-class multi-tool tracking requirements. This dataset also formalized three different *perspectives* of trajectories capturing (1) the life-long intraoperative use of tools, (2) intracorporeal cycle of the tools within the body, and (3) the visibility lifespan of the tools within the camera field of view (FoV) as shown in Fig. 1. Simultaneously tracking tools across these three perspectives is termed *multi-perspective tracking*. The CholecTrack20 dataset provides rich multi-perspective tracking annotations adaptable to diverse surgical needs, however, to date, no deep learning model has been explored on this dataset for automatic tool tracking.

To develop a method for multi-perspective multi-class multi-tool tracking in surgical videos, we first benchmark 10 state-of-the-art

\* Corresponding author at: University of Strasbourg, CAMMA, ICube, CNRS, INSERM, France  
E-mail address: [nwoye@unistra.fr](mailto:nwoye@unistra.fr) (C.I. Nwoye).

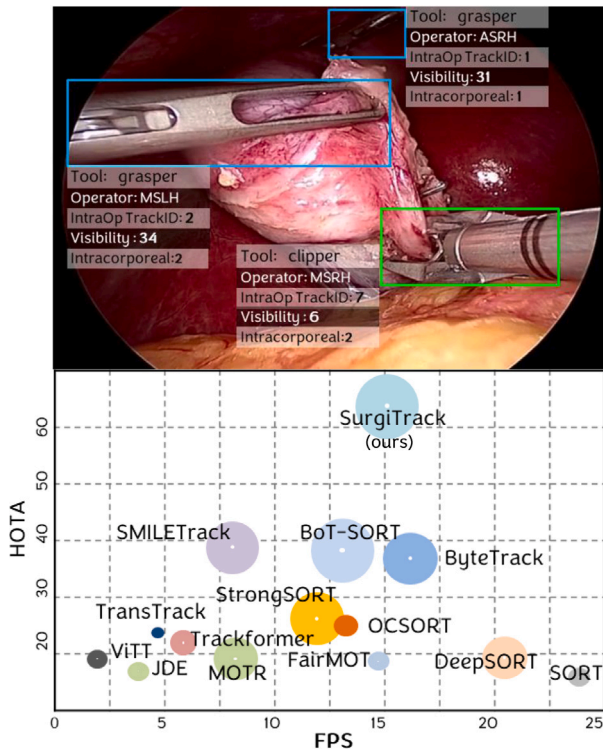


Fig. 1. Surgical tool tracking demonstrating (top) qualitative fine-grained tracking result across multiple tools, classes, and perspectives and (bottom) superior quantitative results compared to the state-of-the-art.

(SOTA) detection methods on the CholecTrack20 dataset and conduct an extensive ablation study on suitable re-ID methods for tracking in the surgical domain. The re-ID module plays a pivotal role of managing the tools identities across time in surgical videos. But, challenges arise from the intricate motion patterns of the tools, frequent occlusions, and the limited field of view within the surgical scene. A particularly daunting task is the re-identification of tools after they have been occluded, moved out of the camera's view, or re-inserted into the surgical field. This complexity is amplified when multiple instances of the same tool class share identical appearance features. Contrary to existing approaches, our preliminary experiments revealed that relying solely on tool appearance cues for track discrimination is sub-optimal especially when distinguishing between instances of the same class. To address this issue, we turn to domain knowledge, specifically the tool's usage pattern and the tool operator's information. The latter criterion, tool operator, refers to the surgeon's hand manipulating the tool and is found to be more accurate than appearance features in distinguishing between instances of the same tool class. However, the tool operators are not directly observable in endoscopic images, making their automatic prediction a challenging endeavor.

Inspired by our findings, we propose *SurgiTrack*, a novel deep learning approach for surgical tool tracking. *SurgiTrack* casts the tool operators as the approximation of the tools' originating directions. It then employs an Attention Mechanism to encode tool motion direction, effectively emulating the unseen surgeon operator's hands or trocar placements for tool re-identification. Our model design allows for an alternative self-supervision of the direction estimator with a comparable performance to its supervised counterpart. This technique ensures that our method can be explored on surgical datasets where the tool operator labels are unavailable. Finally, to account for the multi-perspective nature of tool trajectories, our network associate tracks using a *harmonizing* bipartite matching graph algorithm, which, aside from the usual linear assignment, resolves identity conflicts across

the track perspectives and improves the accuracy of the track ID re-assignment in general.

To summarize, our contributions with this work include the formalization of multi-perspective tool tracking modeling and benchmarking of state of the art methods on the CholecTrack20 dataset. It also includes the development of the *SurgiTrack* model which relies on self-supervised attention-based motion direction estimation and harmonizing bipartite graph matching for tool tracking. Finally, we provided an extensive evaluation of tool tracking on the different trajectory perspectives, at varying video frame rates, and under various visual challenges such as bleeding, smokes, occlusion, etc. These contributions collectively advance the state of the art in surgical tool tracking, opening doors to enhanced computer-assisted surgical systems and AI interventions in the field.

## 2. Related work

Despite the many potentials of tool tracking, challenges persist, prompting our research into more robust and efficient tracking techniques. In the domain of surgical tool tracking, an array of tracking modalities (Fried et al., 1997; Chmarr et al., 2007; Song et al., 2023; Speidel et al., 2008; Behrens et al., 2011; Reiter et al., 2012; Sznitman et al., 2012), ranging from electromagnetic to optical and image-based tracking, has been explored to enhance precision and efficiency. While robotic kinematics gains popularity for its accuracy in instrument localization, image-based tracking stands out for its non-invasiveness, applicability, generalizability, and seamless alignment with the surgeon's visuals. Despite its merits, image-based tracking faces challenges such as occlusion, deformation, and lighting issues. Our work aligns with the trajectory of existing deep learning research (Bouget et al., 2017; Lee et al., 2019; Nwoye et al., 2019; Zhao et al., 2019b,a; Robu et al., 2021; Nwoye, 2021; Fathollahi et al., 2022; Wang et al., 2022; Rueckert et al., 2023), contributing to the ongoing refinement of image-based tracking, positioning it as a pivotal frontier in revolutionizing surgical assistance.

The dynamic landscape of tool tracking extends beyond single-object tracking (SOT) (Danelljan et al., 2017; Zhao et al., 2019b) to encompass multi-object tracking (MOT) (Wojke et al., 2017; Wang et al., 2020; Zhang et al., 2021; Robu et al., 2021) and multi-class tracking (MCT) (Nwoye et al., 2019; Nwoye, 2021). For diverse surgical tools, our research lies within the scope of multi-class multi-object tracking (MCMOT) (Lee et al., 2016; Jo et al., 2017; Zhu et al., 2021; Du et al., 2021), reflecting its potential for real-world surgical applications (Pedrett et al., 2023; Xu et al., 2023, 2022; Dutkiewicz et al., 2005; Richa et al., 2011; Martin-Gomez et al., 2023). In the pursuit of robust tracking methodologies, the choice between tracking paradigms, such as tracking by detection (Zhang et al., 2021; Wang et al., 2020; Aharon et al., 2022; Zhang et al., 2022; Wang et al., 2023b), regression (Bergmann et al., 2019), attention (Sun et al., 2020; Zeng et al., 2022; Meinhardt et al., 2022; Chu et al., 2023), segmentation (Lee et al., 2019) or fusion (Reiter et al., 2012), guides the selection based on context and task suitability. Our research adopts a tracking-by-detection paradigm for its simplicity in achieving surgical tool tracking precision.

Amidst the evolution of tracking methods from rule-based and marker-based approaches (Ma et al., 2021; Huang et al., 2020; Cartucho et al., 2022) to markerless techniques (Reiter et al., 2012; Ye et al., 2016) and machine learning methods (Rieke et al., 2015; Zhao et al., 2019b; Alsheikhali et al., 2015; Pezzementi et al., 2009; Speidel et al., 2008; Dockter et al., 2014), our work is positioned at the forefront, leveraging deep learning advancements like Siamese networks and attention-based networks (Sun et al., 2020; Meinhardt et al., 2022). Notably, our focus on direction features, capturing the tool handling direction by surgeon operators, sets our approach apart, offering a more suitable discriminator for refined tool tracking accuracy in scenarios where appearance (Reiter et al., 2014) or similarity (Wang




Fig. 2. Overview of CholecTrack20 dataset showing localization, tracking, and associated labels (Nwoye et al., 2023).

Table 1

Summary of the CholecTrack20 dataset statistics (Nwoye et al., 2023).

Details		Average track length/number of tracks per perspective				Surgical visual challenges	
Params	Counts	Tools	Intraoperative	Intracorporeal	Visibility	Events	Counts
Videos	20	Grasper	582.0/59	209.4/164	25.7/1334	Occlusion	22 958
Total duration	14h 3m 1s	Bipolar	118.0/20	30.6/78	24.6/96	Bleeding	18 698
Tool categories	7	Hook	1063.7/20	181.8/118	87.5/244	Smoke	2388
Phase categories	7	Scissors	44.6/20	27.0/34	23.4/46	Crowded	6737
Operator categories	4	Clipper	71.7/20	36.8/40	31.9/46	Blurring	299
Labeled frames (1 FPS)	35K	Irrigator	90.8/20	25.9/56	13.4/108	Light reflection	74
Total frames (25 FPS)	1.3M	Spec.Bag	175.2/20	166.8/20	44.9/78	Fouled lens	2196
Bounding boxes	65.2K	Total	2146.0/179	678.3/510	251.4/1952	Poor camera coverage	357

et al., 2023b) features may prove less robust. Contrary to the joint detection and embedding (JDE) setting (Zhang et al., 2021; Wang et al., 2020), adopting an online separate detection and tracking (SDE) approach (Bewley et al., 2016; Maggiolino et al., 2023; Du et al., 2023; Aharon et al., 2022; Du et al., 2021; Wang et al., 2023b), where detection precedes tracking, aligns with the essence of an independently trained tracker that assumes prior detection results. Anchored in dataset like CholecTrack20 (Nwoye et al., 2023), which offers both multi-perspective trajectories and MCMOT requirements in endoscopic videos, our research aligns with the ongoing efforts (Nwoye, 2021; Robu et al., 2021; Fathollahi et al., 2022; Rueckert et al., 2023) to address the complexities of surgical tool tracking, contributing to the broader landscape of advancements in this critical domain.

### 3. Dataset

Our research is conducted on CholecTrack20 (Nwoye et al., 2023) dataset consisting of 20 videos of laparoscopic procedures that have been fully annotated with detailed labels for multi-class multi-tool tracking.<sup>1</sup> The dataset, illustrated in Fig. 2, provides track identities across 3 perspectives of track definition: (1) visibility trajectory of a tool within the camera scope, (2) intracorporeal trajectory of a tool while within a patient's body, and (3) life long intraoperative trajectory of a tool. Intraoperative tracking not only re-identifies tools out of camera view (OOCV) as done in intracorporeal tracking but also maintains their trajectory when out of body (OOB). In the CholecTrack20 dataset, OOB is detected/annotated either by visually observing the tool exit the trocar, inferring from another tool entering through the

same trocar, or noting that the initial tool releases its grasp while out of camera focus. The dataset also provides detailed labels for each tool such as spatial bounding box coordinates, class identity, operator identity, phase identity, frame visual conditions such as occlusion, bleeding, and presence of smoke statuses, among others. The annotated tool categories are grasper, bipolar, hook, scissors, clipper, irrigator and specimen bag. The annotated tool operators are main surgeon left hand (MSLH), main surgeon right hand (MSRH), assistant surgeon right hand (ASRH) and null operator (NULL). The annotations are provided at 1 frame per second (FPS) consisting of 35 K frames and 65 K instance tool labels. Raw videos, recorded at 25 FPS, are provided for inference. The dataset is available on Synapse (Nwoye et al., 2024). A summary of the CholecTrack20 dataset statistics showing the size of key elements in the dataset, number of challenging scenarios, average track duration and number of tracks for each tool category per perspective is presented in Table 1. More information on the dataset can be found in Nwoye et al. (2023).

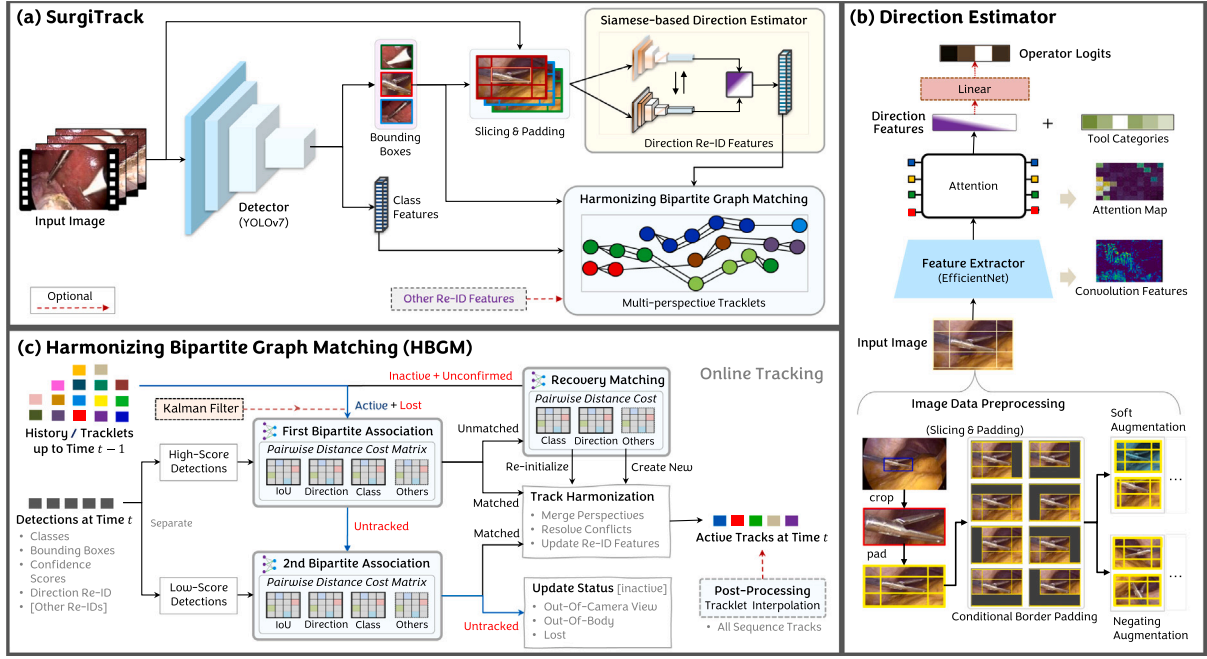
### 4. Methods

We present *SurgiTrack*, a deep learning method for surgical tool tracking based on tool direction of motion features. *SurgiTrack* is designed as a multi-class multi-object tracking (MCMOT) model capable of tracking tools jointly across multiple trajectory perspectives, namely visibility, intracorporeal, and intraoperative. The motivation to track beyond camera's field of view is to offer more flexible trajectories that ensure continuous and reliable identification of surgical tools, tailored to the complex dynamics of a surgical scene, preventing errors and maintaining safety even when tools temporarily move out of view.

The architecture of our proposed tracking model is conceptually divided into the main components of object tracking: spatial detection and data association, with the later further split into re-identification feature modeling and track identity matching, as illustrated in Fig. 3(a).

<sup>1</sup> Dataset will be released before this paper publication.





**Fig. 3.** Overview of our proposed tool tracking model showing: (a) full architecture of SurgiTrack and its major component modules. One of the which is the YOLO-based detector. The other is the Siamese-based surgical tool direction estimator — full architectural detail in (b) which also shows an optional head for surgeon operator classification. The last component of the SurgiTrack is the harmonizing bipartite graph matching (HBGM) algorithm for tool track identity association under multiple perspectives of tool trajectories: visibility, intracorporeal, and intraoperative — full pipeline in (c).

#### 4.1. Multi-perspective tool track formalization

Given a video dataset  $D = \{S_1, S_2, \dots, S_N\}$ , where each sequence  $S_i = \{f_1, f_2, \dots, f_T\}$  consists of frames  $f_i$  containing tools represented by bounding box locations  $B_i = [B_1^i, B_2^i, \dots, B_M^i]$  and associated classes  $C_i = [C_1^i, C_2^i, \dots, C_M^i]$ , the objective is to accurately track tool identities over time across multiple perspectives. For each perspective  $p \in \{1, 2, 3\}$ , we define an association matrix  $A^{(p)}(t)$ , where  $A_{i,j}^{(p)}(t) = 1$  indicates that the  $i$ th detected tool in frame  $f_t$  is associated with the  $j$ th detected tool in frame  $f_{t+1}$ , and  $A_{i,j}^{(p)}(t) = 0$  otherwise. The final tracking solution involves harmonizing these perspective-specific matrices  $\{A^{(1)}(t), A^{(2)}(t), A^{(3)}(t)\}$  to ensure consistent and accurate tool tracking.

#### 4.2. Spatial detection

The spatial detection module is responsible for detecting the tools in each frame of a laparoscopic video. The detection output  $D_t = (B_t^i, C_t^i)$  of frame  $f_t$  is a pair of bounding box coordinates and class identities, with  $B_t^i = [X_t^i, Y_t^i, W_t^i, H_t^i]$  containing the center coordinates  $(X_t^i, Y_t^i)$  and the spatial sizes  $(W_t^i, H_t^i)$  of the boxes and  $C_t^i \in \{0, 1, \dots, 6\}$  represents the class identity of the tools.

We use YOLOv7 (Wang et al., 2023a), a single shot object detector renowned for balancing accuracy and real-time inference, as the backbone for our detection module. Pretrained on the Crowdhuman (CH) (Shao et al., 2018) and MOT20 datasets (Dendorfer et al., 2020), then finetune on the CholecTrack20 data, YOLOv7 produces tool class identities  $C = [0, 1, \dots, 6]$  and bounding boxes  $B = [X, Y, W, H]$ , followed by Non-Maximum Suppression (NMS) at a 0.3 threshold to filter out redundant detections. The predicted bounding boxes are then used to crop out the tool regions from the RGB image for the Re-ID phase.

#### 4.3. Re-identification by direction estimation

Re-identification of tool instances is essential for managing their identities across time. It involves learning representations that can

uniquely identify tool instances across frames. In a video frame where multiple tools can share identical appearance, given that every tool is uniquely tied to a tool operator which remains consistent for their distinct trajectory, automatic estimation of these operators would be helpful in overcoming the challenges of tracking surgical tools in videos. The tools operators, however, are not observable from the captured images. We argue that the operator of a tool can be discerned from the originating direction of the tools termed the *fro-direction*. We therefore propose a 3-step process of (1) data preprocessing, (2) direction estimation, and (3) operator prediction, as shown in Fig. 3(b), to indirectly infer the hidden operators from a given image by learning tool direction.

##### 4.3.1. Data preprocessing

We observed that the tool shaft, while not part of the bounding box annotations, more consistently points toward the *fro-direction* than the tip of the tool, which otherwise aligns with the *to-direction*. However, there is a limitation — the tool shaft may not always be visible in the images, especially when the tool is positioned near the corners of the frame. Our preprocessing technique of *image slicing and padding*, shown in Fig. 3(b), addresses this issue. Here, we crop the images using the tool bounding boxes and pad the cropped images with a percentage of the surrounding neighborhood pixels. Groundtruth and predicted bounding boxes are used at training and testing times respectively. Where a bounding box is at the border of the frame, zero filled padding is applied. This approach helps to reveal the presence or absence of tool shafts.

##### 4.3.2. Direction estimation

We model the estimation of the direction  $\theta$  of a surgical tool, given its bounding box location  $B$ , using a Bayesian-inspired framework. Specifically, the relationship between  $\theta$  and  $B$  follows Bayes' rule as shown in Eq. (1):

$$P(\theta|B) = \frac{P(B|\theta) \cdot P(\theta)}{P(B)} \quad (1)$$

Here,  $P(B|\theta)$  represents the likelihood of the tool's bounding box given its direction,  $P(\theta)$  is the prior probability of the direction, and  $P(B)$

is the evidence, representing the marginal probability of the bounding box across all possible directions.

We approximate this framework using the Attention Mechanism A in Eq. (2), which captures dependencies between bounding box features and direction:

$$\mathbf{A}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{KQ}^T}{\sqrt{d_K}}\right) \mathbf{V} \quad (2)$$

The Attention Mechanism is a common approach in neural networks for learning dependencies between variables (Bahdanau, 2014). In this formulation, the likelihood  $P(B|\theta)$  is approximated by the attention alignment scores  $\text{Softmax}(\mathbf{KQ}^T/\sqrt{d_K})$ , which measure the relevance of each direction based on the bounding box features, i.e., how much attention (weight) each trocar direction should get based on the current tool location. The prior probability  $P(\theta)$  is represented by the value vector  $\mathbf{V}$ , which learns a distribution over possible directions independent of the bounding box features. The evidence  $P(B)$  serves as a normalization term, implicitly handled through the *Softmax* operation. The query  $\mathbf{Q}$ , key  $\mathbf{K}$ , and value  $\mathbf{V}$  are learned embeddings that encode direction vectors, bounding box features, and direction priors, respectively, from the input images using convolutional neural network (CNN) layers. Specifically, we use EfficientNet-b0 (Tan and Le, 2019) to encode these features.

The Attention Mechanism A effectively approximates  $P(\theta|B)$  by learning the alignment between bounding box features and directional priors within a differentiable framework. Unlike classical Bayesian methods where priors and likelihoods are explicitly defined, this approach allows the network to learn these distributions directly from the data, providing flexibility for modeling complex spatial dependencies. By grounding the formulation in Bayes' rule and leveraging the interpretability of attention mechanisms, we provide a principled approach to estimating tool direction while maintaining compatibility with modern deep learning architectures.

#### 4.3.3. Operator estimation

The tool operator  $O$  is estimated as the conditional probability  $P(O | [\mathbf{A}, \mathbf{C}])$ , where  $\mathbf{A}$  represents the attention-based direction features (derived from Eq. (2)), and  $\mathbf{C}$  is the tool category. This estimation is modeled using a fully connected layer, with the linear transformation defined as:

$$\mathbf{y} = \mathbf{w} \cdot [\mathbf{A}, \mathbf{C}] + \mathbf{b} \quad (3)$$

where,  $\mathbf{w}$  and  $\mathbf{b}$  are the weight and bias parameters, and the input is the concatenation of  $\mathbf{A}$  and  $\mathbf{C}$ . To obtain the final probability distribution, a SoftMax function is applied to  $\mathbf{y}$ :

$$P(O | [\mathbf{A}, \mathbf{C}]) = \text{SoftMax}(\mathbf{y}) \quad (4)$$

This process generates the probability of each operator  $O$  given the tool direction and category, enabling operator classification.

#### 4.3.4. Supervision

Our Direction Estimation Network can be trained through multiple paradigms, depending on the availability of labels. For *full supervision*, the model is trained on and to predict the operator ID labels, employing a weighted cross-entropy loss for class balancing as in Eq. (5), where  $y$  and  $\hat{y}$  are respectively the groundtruth and predicted labels for the operators,  $\sigma$  is a sigmoid function, and  $w$  is a weight vector for class balancing.

$$L = -1 \left( y \log(\sigma(\hat{y}))w + (1 - y) \log(1 - \sigma(\hat{y})) \right) \quad (5)$$

In *weak- and self-supervised* learning, the model adopts a Siamese-style approach, contrasting negative or dissimilar pairs and pulling positive or similar pairs closer using the margin-based loss in Eq. (6), where  $y$  is 0 for negative pairs and 1 for positive pairs,  $d$  is the

Euclidean distance between two pairs, and  $m$  is the margin of distance tolerance, set to 0.5 in our experiment.

$$L = d^2 y + \max(m - d, 0)^2 (1 - y) \quad (6)$$

In the weak supervision scheme, we leverage the knowledge that tools with different track IDs have different operator categories, helping segregate positive and negative pairs. For the self-supervised setting, we group pairs based on the tool's *direction assumption*: in a frame where there are multiple tools, they are considered opposing pairs and their positive pairs are generated by a soft augmentation (e.g. scale, perspective, slight rotation  $\theta < 10^\circ$ , mixed). Conversely, in frames with a single tool, negative pairs are generated by rotating the tool to a different direction (usually  $80^\circ \leq \theta \leq 100^\circ$  and  $170^\circ \leq \theta \leq 190^\circ$ ), and positive pairs are formed through soft augmentation.

#### 4.4. Identity association using multi-perspective matching

The role of the identity association module is to link tool detections across frames, creating coherent tracks for surgical tools. It is a critical step in understanding how these tools move and interact during procedures. To tackle this challenge effectively, we introduce a novel approach called *Harmonizing Bipartite Graph Matching* (HBGM). This method is designed to handle the diverse perspectives of surgical tool trajectories without conflicts, ensuring that every tool is tracked accurately throughout the surgery. HBGM builds upon the Bipartite graph matching (BGM) algorithm (Kuhn, 1955), which is a powerful tool in its own right. What sets HBGM apart is its ability to manage multiple track states for each tool instance. These states include:

1. *New*: Assigned at track initialization.
2. *Active*: Updated when a detection matches a track.
3. *Lost*: Assigned to an active track with no matching detection at current time. The "Lost" state has a 5-second time-threshold within which the track can be re-activated for all perspectives.
4. *Out of Camera View (OOCV)*: Assigned if a track remains "Lost" for more than the time-threshold. This state marks the end of a visibility trajectory. An "OOCV" track can only be re-activated for intraoperative and intracorporeal perspectives.
5. *Out of Body (OOB)*: Assigned if a tool track is judged to have left the abdomen. This is determined when a track fails a class-direction consistency check (described in Section 4.4.2). The OOB state marks the end of an intracorporeal trajectory. Such track can only be re-activated for intraoperative perspective.
6. *Removed*: Assigned if a track is initialized erroneously (e.g., after the "New" state, it never progresses to "Active" for a prolonged time) or not reassigned for a long period. Removed tracks are excluded from history against a potential match.

The stages of HBGM are illustrated in Fig. 3(c) and discussed as follows:

##### 4.4.1. Pairwise matching

To match tools across frames, we use pairwise distances between predicted motion directions ( $\theta$ ) and detection boxes ( $B$ ), while ensuring class ( $C$ ) consistency. For each predicted tool  $p_i^t = \{B_i^t, C_i^t, \theta_i^t\}$  in frame  $f_t$ , and each active track  $q_j^{t-1} = \{B_j^{t-1}, C_j^{t-1}, \theta_j^{t-1}, \text{IDs}\}$  in the history trajectories up to frame  $f_{t-1}$ , the distance metric  $\text{Dist}_{i,j}$  is defined as:

$$\begin{aligned} \text{Dist}_{i,j} = & \alpha \cdot \text{Dist}_{\text{spatial}}(B_i^t, B_j^{t+1}) + \\ & \beta \cdot \text{Dist}_{\text{direction}}(\theta_i^t, \theta_j^{t+1}) + \\ & \gamma \cdot \text{Dist}_{\text{class}}(C_i^t, C_j^{t+1}), \end{aligned} \quad (7)$$

where

$$\text{Dist}_{\text{spatial}}(B_i^t, B_j^{t+1}) = 1 - \frac{|B_i^t \cap B_j^{t+1}|}{|B_i^t \cup B_j^{t+1}|}, \quad (8)$$

is the Intersection over Union (IoU) similarity between the bounding boxes  $B_i^t$  and  $B_j^{t+1}$ ,

$$\text{Dist}_{\text{direction}}(\theta_i^t, \theta_j^{t+1}) = 1 - \frac{\theta_i^t \cdot \theta_j^{t+1}}{\|\theta_i^t\| \cdot \|\theta_j^{t+1}\|}, \quad (9)$$

is the angular difference between motion directions, and

$$\text{Dist}_{\text{class}}(C_i^t, C_j^{t+1}) = \begin{cases} 1 - S & \text{if } C_i^t = C_j^{t+1} \\ \tau & \text{if } C_i^t \neq C_j^{t+1}, \end{cases} \quad (10)$$

ensures that tools of different classes are not matched ( $S$  is the class probability score of the tool and  $\tau$  is set above matching threshold). Here,  $\alpha$ ,  $\beta$ , and  $\gamma$  are weights that balance the influence of spatial, motion, and class distances, respectively. A bipartite matching algorithm is then used to minimize  $\text{Dist}_{i,j}$  and establish correspondences. Optionally, a second bipartite association is performed for low-confidence detections, following the method proposed by Zhang et al. (2022).

#### 4.4.2. Track harmonization

The magic of HBGM happens during the harmonization stage, where it skillfully handles track ID assignment ensuring consistent tool tracking across multiple perspectives — namely intraoperative, intracorporeal, and visibility. The harmonization process relies on a structured approach involving (a) class-direction consistency checks, (b) hierarchical ID assignment, and (c) referential integrity to effectively manage tool trajectories under different states and handle potential conflicts in track identification and state updates.

**A. Class-Direction Consistency Check:** HBGM validates that a track maintains a joint matching of class ( $C$ ) and direction ( $\theta$ ) features across time to remain active. This check is specifically for managing the intracorporeal trajectory. A change in the class feature (or class ID) while the direction feature remains unchanged indicates that a new tool  $j$  has been introduced through the same trocar as previous tool  $i$ , marking the end of the tool  $i$  intracorporeal trajectory and the start of a new one:

$$\text{If } \theta_i^t = \theta_j^{t-1} \text{ and } C_i^t \neq C_j^{t-1}, \text{ then State}_i^{t-1} = \text{OOB}. \quad (11)$$

This approach enforces a one-to-one correspondence between track direction and class features for more stable and accurate tracking.

**B. Hierarchical ID Initialization and State Updates:** HBGM employs a hierarchical approach to manage new ID assignments and state updates effectively across perspectives. At the topmost level is the intraoperative trajectory ( $p = 3$ ) followed by the intracorporeal trajectory ( $p = 2$ ), and lastly the visibility trajectory ( $p = 1$ ).

1. **ID Initialization:** A new ID in a higher perspective ( $p = 3$ ) triggers new IDs in all lower perspectives ( $p = 2, p = 1$ ). Similarly, new IDs initialized in  $p = 2$  are also initialized in  $p = 1$ .
2. **State Updates:** State changes in lower perspectives are propagated in higher ones. For example, “Lost” in  $p = 1$  implies “Lost” in  $p = 2$  and  $p = 3$ . “OOBV” in  $p = 2$  implies “OOBV” in  $p = 3$ . Hierarchical state update is excluded for “New” state, which is treated as ID initialization.

**C. Referential Integrity:** Once a tool’s track is established in one perspective, every changes in its state is concurrently reflected across all other perspectives. HBGM improves consistent track assignment for unmatched detections and untracked tracks by cross-referencing of historical data across perspectives. This minimizes ID switches and facilitates track recovery after occlusion, OOCV or OOB. Also, persistent trajectory patterns in  $p = 3$  is relied on to dynamically adjust thresholds for the inactive tracks, aiding in robust recovery.

**Table 2**

Experiment and hyperparameter settings.

Pre-training	MOT20, CH	Learning rate	3e−4
Pre-processing	Crop, Resize	Batch size	32
Post-processing	NMS	Epochs	132
Data augmentation	Soft, Negative	Optimizer	Adam
Detector input shape	542 × 655	Weight decay	1e−5
Re-ID input shape	224 × 224	Model params	123M
Train/val/test splits	10/2/8 videos	Lr decay schedule	Plateaux
Detector loss func.	BCE	GPU resources	4 × V100
FSL Re-ID loss	BCE	Exp. training time	48 h
WSL Re-ID loss	Cosine CL	Baseline training time	139 h
SSL Re-ID loss	Contrastive	Ablation training time	120 h

#### 4.4.3. Track recovery

Before initializing new tracks, the algorithm attempts to recover unmatched detections at time  $t$  by pairing them with inactive tracks as at time  $t - 1$  using only directional ( $\theta$ ) and categorical ( $C$ ) features. For each unmatched detection  $d_i^t \in \mathbf{D}_{\text{unmatched}}^t$  and inactive track  $d_j \in \mathbf{D}_{\text{inactive}}^{t-1}$ , a recovery score  $S_{ij} = \alpha \cdot \text{Dist}_{\text{direction}}(\theta_i^t, \theta_j) + \beta \cdot \mathbb{1}(C_i^t \neq C_j)$  is computed, where  $\text{Dist}_{\text{direction}}(\theta_i^t, \theta_j)$  measures directional similarity and  $\mathbb{1}(C_i^t \neq C_j)$  indicates categorical mismatch. The unmatched detection  $d_i^t$  is assigned to the inactive track  $d_j$  with the lowest score, provided it falls below a predefined threshold  $\tau_{\text{rec}}$ . If no valid match is found, a new track ID is initialized. This approach leverages feature consistency to ensure seamless recovery and reduces the chance of ID switching or track loss.

#### 4.4.4. Versatility

Optionally, HBGM can seamlessly integrate other re-identification features, such as the Kalman filter (Kalman, 1960), appearance features (Zhang et al., 2021), camera motion compensation (Aharon et al., 2022), and similarity features (Wang et al., 2023b).

The Kalman filter (Kalman, 1960) enhances the spatial distance metric term ( $\text{Dist}_{\text{spatial}}$ ) in Eq. (7) by predicting the future positions of track bounding boxes based on their previous states. The Euclidean distance between these predicted positions and the current bounding box centers is then calculated, yielding a more robust spatial cost component. For appearance and other re-ID features, a corresponding distance cost is computed similarly to the direction features ( $\text{Dist}_{\text{direction}}$ ). Each re-ID feature contributes its distance term, which is weighted and added to the overall distance metric  $\text{Dist}_{i,j}$  in Eq. (7) before applying bipartite matching for track assignment.

Overall, our approach ensures accurate track ID reassignment, mitigates identity fragmentation and switches, and provides fine-grained track details and a rich trajectory history.

#### 4.5. Baseline method

To demonstrate the efficacy of our sophisticated modeling, we design *knowledge-based (KB) baseline*, integrating predefined surgical states and priors into existing trackers. In the context of laparoscopic cholecystectomy, approximately 75% of the utilized tools in a procedure exhibit unique categories and a known instance count, applicable solely within the intraoperative trajectory perspective. The KB baseline is a straightforward, non-trainable algorithm that incorporates predefined constraints:

1. **Inter-Class Switch Constraint:** This utilizes the tool class features to disallow identity switch across tool categories with inter-class bounding box overlap, mostly observed in crowded scenes and tool occlusion scenarios.
2. **Maximum Instance Constraint:** To mitigate identity fragmentation, a new track assignment is allowed only if the assigned trajectories per tool category remain within the allowable maximum track instances otherwise the matching threshold is incrementally lowered until the candidate tool is matched to a

**Table 3**

Surgeon operator prediction based on direction feature embeddings for tool track re-identification — results showing the mean and per-class average precision (% AP) and embedding temporal consistency (% Accuracy) from time  $t - k$  to  $t$ . (Grasper only).

Model	Full supervision								Weak supervision				Self supervision			
	MSLH	ASRH	MSRH	mAP	$T'_{t-1}$	$T'_{t-5}$	$T'_{t-25}$	$T'_{t_0}$	$T'_{t-1}$	$T'_{t-5}$	$T'_{t-25}$	$T'_{t_0}$	$T'_{t-1}$	$T'_{t-5}$	$T'_{t-25}$	$T'_{t_0}$
Siamese baseline	92.3	81.5	11.2	61.7	80.0	78.1	74.7	68.1	<b>64.6</b>	<b>62.1</b>	58.2	<b>51.1</b>	62.1	61.0	58.6	47.8
ViT	92.1	77.7	10.2	60.0	78.7	76.7	73.2	67.3	63.0	58.7	53.4	50.3	64.8	62.1	57.7	51.0
CrossViT	91.7	79.6	9.7	60.3	78.5	76.8	73.9	69.8	61.7	60.0	57.5	50.3	61.8	60.7	58.2	47.4
ResNet-18	95.5	87.1	36.6	73.0	81.3	79.2	76.3	72.8	62.1	60.8	57.7	46.5	87.6	86.7	85.4	83.5
EfficientNet-B0	<b>97.5</b>	92.1	32.5	74.0	86.0	84.5	82.7	81.3	<b>64.6</b>	61.5	55.3	42.8	89.7	88.7	87.3	85.1
SMILETrack	89.8	75.7	13.6	59.7	78.5	76.9	73.7	62.3	62.0	60.0	54.5	50.6	61.7	60.4	56.7	46.1
Proposed estimator	<b>97.5</b>	<b>93.0</b>	<b>53.0</b>	<b>81.2</b>	<b>91.0</b>	<b>89.0</b>	<b>88.6</b>	<b>87.4</b>	63.6	61.8	<b>58.3</b>	47.5	<b>91.6</b>	<b>90.7</b>	<b>89.9</b>	<b>88.4</b>

probable existing trajectory. Attempts to extend the application of this constraint beyond intraoperative tracking, with non-deterministic instance count, yield no meaningful results.

Aside from the KB baseline, we benchmark over 20 SOTA Deep Learning methodologies on CholecTrack20 dataset for surgical tool detection, re-identification, and tracking. We adapt some of the SOTA tracking models for MCMOT to suit the surgical context of our study.

## 5. Experiments

### 5.1. Implementation details

The proposed model is implemented in PyTorch; implementation details and hyperparameters are tabulated in Table 2. Depending on code availability, baseline models are outsourced from public repositories or re-implemented. We follow the official splits of the CholecTrack20 (Nwoye et al., 2023) dataset to set 10, 2, and 8 videos for training, validation, and testing respectively.

### 5.2. Evaluation metrics

We assess the quality of the direction embeddings for track re-identification through a consistency accuracy metric: here, a predicted tool's direction is correct if the pairwise difference between its embeddings at time  $t$  and  $t - 1$  is less than a predefined threshold (usually  $\theta = 0.5$ ). We utilize extended time intervals  $t - k$ , comparing embeddings consistency up till the past  $k$  frames,  $k \in [1, 5, 25, \text{start}]$ .

We follow the recommended protocol and metrics for CholecTrack20 (Nwoye et al., 2023) tool tracking evaluation. Here, we evaluate different aspects of the tracking performance using the standard metrics of higher-order tracking accuracy (HOTA) (Luiten et al., 2021) (including its constituents: Localization Accuracy, LocA, Detection Accuracy, DetA, and Association Accuracy, AssA), the CLEAR metrics (Bernardin and Stiefelhagen, 2008) of multi-object tracking accuracy (MOTA), multi-object tracking precision (MOTP), mostly tracked (MT), mostly lost (ML), partially tracked (PT), identity switch (IDSW), track fragmentation (Frag). We also assess the identity F1 score (IDF1) (Ristani et al., 2016), counting metrics (number of detections #Dets and identities #IDs), and tracking speed in frame per second (FPS).

## 6. Results and discussion

First, our base detector, YOLOv7 (Wang et al., 2023a), yields 80.6%  $AP_{0.5}$  and 56.1%  $AP_{0.5:0.95}$  for tool detection at an inference speed of 20.6 FPS, demonstrating its effectiveness as a detector for our tracking model.

### 6.1. Results of surgeon operator prediction

We measure the quality of the direction features used for track re-identification in this task. The proposed Estimator, built on the EfficientNet-B0 (Tan and Le, 2019) backbone, demonstrates remarkable performance in surgical tool re-identification across different supervi-

sion settings as evidenced by the results presented in Table 3. Leveraging EfficientNet-B0's known efficiency and speed, it outperforms several strong baselines: Siamese baseline, ResNet (He et al., 2016), ViT (Dosovitskiy et al., 2020), CrossViT (Chen et al., 2021), etc., in capturing essential features for tool re-identification and achieving the highest mean Average Precision (mAP) of 81.2% under supervised setting. The inclusion of an attention head enhances its ability to learn direction features, which is crucial for distinguishing between similar tool instances, thereby outperforming the baseline EfficientNet-B0 (+4.2% mAP). Category-wise, the main surgeon's right hand (MSRH), which is the busiest and handles most of the tools, exhibits the greatest detection difficulty with a 53% AP.

In the self-supervised setting, the proposed Estimator showcases its versatility by consistently achieving high re-identification accuracy ( $\geq 88\%$ ) over various time intervals, including longer time differences up to the start of the video (e.g.,  $T'_{t-25}$  and  $T'_{t_0}$ ). This demonstrates that the learned direction-aware features effectively maintain consistency, even when dealing with challenging long-term tracking scenarios; thanks to the innovative image preprocessing technique of “*image slicing and padding*” that effectively addresses the visibility issue of tool shafts in images, enhancing the model's ability to capture directional information. The direction features align closely with the tool operator's hand direction from the trocar port, emphasizing their relevance as robust re-identification features for surgical tool tracking. This analysis focus on grasper because it is the only tool with multiple instance when considering intraoperative trajectory. We, however, observe a slightly inferior performance when models are weakly-supervised on track ID labels, which is the same supervisory signal for conventional appearance re-ID models (Zhang et al., 2022; Aharon et al., 2022; Wang et al., 2023b).

### 6.2. Ablation study on Re-ID features

The ablation study in Table 4 explores the impact of track re-identification (Re-ID) features on multi-object tracking, categorizing them into parametric and non-parametric groups. The non-parametric features are bounding box location (IoU), low-confidence detection (Byte) (Zhang et al., 2022), motion prediction using Kalman filter (KF) (Kalman, 1960), camera-motion compensation (CMC) using BoT-SORT-ReID (Aharon et al., 2022), and multi-class (MC) identity labels. The parametric features are appearance features (AF) (Zhang et al., 2021), similarity features (SF) (Wang et al., 2023b), and the proposed direction features (DF) — Using the SSL paradigm.

In the non-parametric settings, IoU, BYTE, and KF show incremental contributions, with IoU + BYTE + MC achieving the highest HOTA (16.8%). Parametric features like AF, SF, and DF exhibit significant individual capabilities, with DF standing out at 59.5% HOTA. Combinations of AF + SF (32.4%) and AF + DF (62.6%) highlight synergies between appearance and directional features. Jointly, IoU + MC + DF excels (60.2%), and further additions (BYTE, AF) refine results, culminating in the optimal IoU + AF + SF + DF (61.5%) configuration. This underscores the nuanced interplay of Re-ID features, with DF proving effective independently, synergistically with AF and SF, and being judiciously combinable for enhanced tracking performance. The



**Table 4**

Ablation study on track re-identification features. Using the intraoperative trajectory perspective @ 25 FPS.

	Re-ID features								Tracking results		
	IOU	BYTE	KF	CMC	MC	AF	SF	DF	HOTA↑	MOTA↑	IDF1↑
Non-parametric	✓								14.8	57.3	8.7
	✓	✓							14.7	59.2	8.8
	✓		✓						13.9	<b>62.2</b>	7.7
	✓			✓					13.5	60.7	7.3
	✓				✓				16.6	58.7	9.6
	✓	✓			✓				<b>16.8</b>	<b>57.2</b>	<b>9.8</b>
	✓	✓	✓	✓	✓				13.9	60.5	7.5
Parametric						✓			19.1	51.3	12.1
							✓		25.5	40.9	19.4
						✓	✓	✓	59.5	72.5	65.4
						✓	✓	✓	32.4	73.3	31.9
						✓	✓	✓	<b>62.6</b>	<b>78.8</b>	<b>71.4</b>
						✓	✓	✓	62.2	77.7	70.9
						✓	✓	✓	62.0	78.2	70.3
Joint	✓					✓			22.8	64.8	16.3
	✓						✓		29.1	49.7	25.0
	✓							✓	<b>61.3</b>	78.5	<b>67.8</b>
	✓				✓			✓	60.2	<b>78.8</b>	66.8
	✓	✓			✓			✓	57.9	75.9	63.6
	✓					✓	✓		25.2	66.7	20.7
	✓					✓		✓	60.7	<b>78.9</b>	67.6
	✓						✓	✓	60.6	77.8	68.0
	✓					✓	✓	✓	<b>61.5</b>	78.5	<b>69.3</b>
	✓	✓			✓	✓	✓	✓	57.6	77.3	64.2
	✓	✓	✓	✓	✓	✓	✓	✓	55.5	65.4	62.1

**Table 5**

Ablation study on linear association algorithms and approach for combining multiple re-ID costs.

Algorithm	Re-ID costs ensemble						Results		
	BM	HBM	Min	Avg	w.Avg	Vote	w.Vote	w.A.V	
✓			✓						12.2
✓	✓		✓						23.1
✓		✓							26.8
✓			✓						37.4
✓				✓					79.9
✓					✓				39.8
✓						✓			36.7
✓							✓		74.9
✓								✓	33.2
✓									56.4
✓									<b>80.1</b>
✓									60.0
✓									58.0
✓									74.4
✓									63.1
✓									<b>61.5</b>
✓									78.5
✓									<b>69.3</b>

study emphasizes the versatility of DF, showing its competence with basic geometric features and in comprehensive multi-feature scenarios, providing insights for real-time applications. Incorporating DF, as shown through the ablation experiments in Table 4, usually leads to a more significant improvement compared to incorporating appearance-based features, suggesting that the DF capture more than just visual and spatial cues.

### 6.3. Ablation study on track association algorithm

The comparative analysis between Bipartite Graph Matching (BGM) and Harmonizing Bipartite Graph Matching (HBGM) in Table 5 reveals a substantial enhancement with HBGM, exhibiting a noteworthy increase of +17.1% HOTA, +35.4% MOTA, and +8.5% IDF1. Regarding the use of multiple re-ID features, we explore the location, appearance, similarity, and direction cost matrices and examine several possible ensemble approaches. Contrary to the commonly used minimum (Min) method, which tends to bias predictions toward the least cost matrix, ensemble methods such as averaging (Avg) and voting (Vote) gives higher scores showcasing the importance of consensus. The introduction of weighted averages (w.Avg) and weighted voting (w.Vote) results in a more optimal solution, emphasizing the significance of weighting by the individual strengths of the re-ID features for more robust tracking. Specifically, averaging minimizes tracking error by reducing variances leading to improved MOTA score. Voting improves identity

assignment consensus resulting in higher HOTA and IDF1 scores. Combining the strengths of weighted averaging and voting (w.A.V) balances and improves the tracking performances.

### 6.4. Analysis of multi-tool tracking results

We evaluate our method on the CholecTrack20 dataset, and compare it with the state-of-the-art and baseline methods (Table 6). The existing models, including OCSORT (Maggiolino et al., 2023), TransTrack (Sun et al., 2020), ByteTrack (Zhang et al., 2022), Bot-SORT (Aharon et al., 2022), and SMILETrack (Wang et al., 2023b), exhibit varying degrees of performance across HOTA metrics. While these models show good detection and localization capabilities with relatively high DetA, LocA, and MOTP scores, they struggle with tool track identity association, evident in their low AssA and IDF1 scores. TransTrack is affected the most with a record low HOTA of 7.4%. ByteTrack, Bot-SORT, and SMILETrack, although relatively better, still show room for improvement in HOTA scores ranging from 15.7% to 17.4%. This illustrates that, given the similarity of most tools especially the ones from the same category, relying on location and appearance/similarity features is not sufficient for their correct re-identification across time.

The infusion of surgical knowledge base (KB) into ByteTrack, Bot-SORT, and SMILETrack (resulting in KB variants) helps the identity association and brings about significant improvements, particularly in HOTA scores, ranging from 36.5% to 37.5%. Despite these enhancements, the proposed model, SurgiTrack, surpasses all existing and KB-infused models, achieving an outstanding HOTA score of 67.3% in the fully supervised setting. SurgiTrack excels in individual HOTA metrics, with high DetA (70.8%) and LocA (86.7%), emphasizing its precision in tool detection and localization. The robust AssA score of 64.1% underscores SurgiTrack's proficiency in tool association or re-identification, thanks to the tool direction features.

Comparatively, the SurgiTrack variants trained in weakly supervised (WSL) or self-supervised (SSL) manners also demonstrate strong performance, with SSL showing a slightly higher AssA of 50.6%. Notably, WSL and SSL produce the highest MOTA and lowest ID switches. Being the more competitive variant, the SSL-based SurgiTrack is preferred in scenarios with limited annotated data. Our analysis also cover the counts of number of detections made, number of unique

**Table 6**

Multi-perspective multi-tool tracking results @ 25 FPS.

Model	HOTA metrics				CLEAR metrics					Identity metrics			Count metrics		Speed
Tracker	HOTA↑	DetA↑	LocA↑	AssA↑	MOTA↑	MOTP↑	MT↑	PT↓	ML↓	IDF1↑	IDSW↓	Frag↓	#Dets	#IDs	FPS↑
Intraoperative trajectory (Groundtruth counts: #Dets = 29994, #IDs = 70)															
OCSORT	14.6	52.7	<b>86.7</b>	4.1	49.2	<b>85.0</b>	24	32	14	9.5	2921	2731	21 936	3336	10.2
FairMOT	5.8	25.8	75.9	1.3	5.0	73.9	3	24	43	4.3	4227	1924	15 252	4456	14.2
TransTrack	7.4	31.5	84.4	1.7	4.2	82.9	9	36	25	4.2	4757	<b>1899</b>	21 640	4079	6.7
ByteTrack	15.8	70.6	85.7	3.6	67.0	84.0	54	12	2	9.5	4648	2429	28 440	5383	16.4
Bot-SORT	17.4	70.7	85.4	4.4	69.6	83.7	<b>58</b>	11	<b>1</b>	10.2	3907	2376	29 302	4501	8.7
SMILETrack	15.9	71.0	85.5	3.7	66.4	83.8	55	13	2	9.2	4968	2369	28 821	5761	11.2
ByteTrack (KB)	36.5	70.1	85.4	19.1	74.5	83.7	<b>58</b>	10	2	33.3	2046	2406	<b>29 504</b>	427	16.3
Bot-SORT (KB)	37.4	67.8	85.4	20.7	73.5	83.6	51	12	7	35.5	1638	2199	27 605	315	8.6
SMILETrack (KB)	37.5	65.7	85.6	21.6	71.4	83.8	49	13	8	35.7	1444	2021	26 517	266	11.0
SurgiTrack (FSL)	<b>67.3</b>	70.8	<b>86.7</b>	<b>64.1</b>	76.0	<b>85.0</b>	48	20	2	<b>81.7</b>	1891	2489	27 499	<b>72</b>	15.3
SurgiTrack (WSL)	56.5	<b>71.7</b>	86.5	44.7	<b>78.9</b>	84.8	54	14	2	60.3	<b>1352</b>	2368	28 120	79	15.3
SurgiTrack (SSL)	60.2	<b>71.7</b>	86.6	50.6	78.8	84.8	54	14	2	66.8	1373	2364	28 102	85	15.3
Intracorporeal Trajectory (Groundtruth counts: #Dets = 29994, #IDs = 247)															
OCSORT	23.7	51.4	86.5	11.0	47.1	84.8	115	87	45	18.1	2953	2796	21 797	3526	10.2
FairMOT	7.5	19.7	76.1	2.9	5.4	74.0	19	60	168	6.0	2890	1496	11 287	3962	14.2
TransTrack	13.1	31.5	84.4	5.5	4.6	82.9	80	79	88	8.7	4648	1791	21 640	4079	6.7
ByteTrack	24.7	70.6	85.7	8.7	67.4	84.0	176	48	23	16.9	4515	2290	28 440	5383	16.4
Bot-SORT	27.0	70.7	85.4	10.4	70.0	83.7	<b>188</b>	38	21	18.9	3771	2238	<b>29 300</b>	4501	8.7
SMILETrack	24.9	66.7	85.5	8.9	66.7	83.8	186	39	22	16.9	4868	2232	28 820	5779	11.2
SurgiTrack (FSL)	<b>55.2</b>	70.8	<b>86.6</b>	<b>43.2</b>	74.8	<b>84.9</b>	170	55	22	<b>61.7</b>	2257	2350	27 499	<b>1092</b>	15.3
SurgiTrack (WSL)	34.6	<b>71.7</b>	86.5	16.8	75.7	84.8	184	41	22	28.9	2300	2229	28 120	2257	15.3
SurgiTrack (SSL)	39.4	<b>71.7</b>	86.5	21.8	<b>77.5</b>	84.8	184	42	<b>21</b>	36.2	<b>1761</b>	2225	28 102	1480	15.3
Visibility trajectory (Groundtruth counts: #Dets = 29994, #IDs = 916)															
SORT	17.4	39.5	85.2	7.8	21.4	83.3	139	399	378	13.4	6619	2138	16 595	8844	19.5
OCSORT	37.0	52.6	86.5	26.2	50.2	84.8	300	371	245	35.9	2317	2260	22 197	3587	10.2
FairMOT	15.3	25.0	75.8	9.5	7.1	73.7	58	218	640	14.4	3140	1574	15 338	4875	14.2
TransTrack	19.2	31.6	84.4	11.8	5.8	82.9	224	280	412	16.1	4273	<b>1403</b>	21 640	4079	6.7
ByteTrack	41.5	70.7	85.7	24.8	69.3	84.0	591	217	108	36.8	3930	1704	28 440	5383	16.4
Bot-SORT	44.7	70.8	85.5	28.7	72.0	83.7	<b>638</b>	184	<b>94</b>	41.4	3183	1638	<b>29 300</b>	4505	8.7
SMILETrack	41.3	71.0	85.6	24.4	68.9	83.8	619	192	105	36.5	4227	1641	28 821	5752	11.2
SurgiTrack (FSL)	61.6	70.9	<b>86.6</b>	53.9	75.4	<b>84.9</b>	551	244	121	68.9	2095	1778	27 499	<b>1814</b>	15.3
SurgiTrack (WSL)	58.4	<b>71.7</b>	86.5	47.8	76.8	84.8	601	201	114	62.3	1980	1650	28 120	2678	15.3
SurgiTrack (SSL)	<b>62.8</b>	<b>71.7</b>	86.5	<b>55.3</b>	<b>78.0</b>	84.8	598	202	116	<b>69.4</b>	<b>1546</b>	1648	28 102	2153	15.3

**Table 7**

Class-wise tracking accuracy and Impact of KB algorithm on state of the art models; [Intraoperative Perspective @ 25 FPS].


Method		HOTA	MOTA	IDF1	Class-wise HOTA						
					Grasper	Bipolar	Hook	Scissors	Clipper	Irrigator	spec.bag
Without KB	ByteTrack	15.8	67.0	9.5	13.3	18.0	10.5	16.3	18.3	5.7	17.6
	Bot-SORT	17.4	69.6	10.2	14.3	19.6	11.4	16.9	19.1	6.4	18.7
	SMILETrack	15.9	66.4	9.2	13.0	18.0	10.4	17.0	18.0	5.9	18.8
With KB	ByteTrack	36.5	74.5	33.3	29.1	24.2	41.3	14.6	22.7	12.5	24.1
	Bot-SORT	37.4	73.5	35.5	33.0	28.9	39.8	16.3	27.7	14.6	19.1
	SMILETrack	37.5	71.4	35.7	31.9	30.4	40.8	17.2	29.2	14.5	24.8
Proposed	SurgiTrack (FSL)	<b>67.3</b>	76.0	<b>81.7</b>	<b>60.8</b>	<b>49.3</b>	<b>75.6</b>	<b>26.6</b>	<b>40.9</b>	<b>24.9</b>	<b>38.1</b>
	SurgiTrack (WSL)	56.5	<b>78.9</b>	60.3	36.8	47.9	75.0	25.5	39.4	24.0	37.2
	SurgiTrack (SSL)	60.2	78.8	66.8	45.8	48.0	75.1	25.5	39.5	24.1	37.2

identities assigned, and the tracking speed as illustrated in Table 6. SurgiTrack shows closer to the groundtruth identity counts. These results demonstrate the effectiveness of our proposed method and underscore its potential as an advanced solution for multi-tool tracking in intraoperative scenarios.

### 6.5. Analysis of multi-perspective tracking results

Looking at the different trajectory perspectives, it is shown in Table 6 that visibility tracking is the easiest with most of the existing models showcasing their strengths. This is expected because deep learning models mostly rely on visual cues, which are captured by camera in the visibility track scenario. The SSL variant of SurgiTrack record a landslide top performance scores of 62.8% HOTA, 78.0% MOTA, and 69.4% IDF1 in visibility tracking. The intracorporeal tracking is

the most challenging since the major factors marking the entry and exit of the tools from the body are not readily visible. SurgiTrack however leverages the proposed direction-aware features, which could be linked to the surgeon operators hands and its rich fine-grained history to estimate the out-of-view and out-of-body status of the tools and outperforms all the existing methods with wide margin. The intra-operative trajectory comes in the middle in terms of difficulty. While it may be challenging to ascertain the persistence of a trajectory after re-insertion, the class features are also helpful especially for tools of different categories. The direction features, however, has a better tendency of estimating the persistent identity of different tools of the same class with a +29.8% and +29.9% HOTA higher than similarity and appearance features respectively. Remarkably, SurgiTrack jointly handle the 3 trajectory perspectives unlike the compared methods, where we train separate model for each perspective.



**Fig. 4.** Impact of direction estimation in tracking surgical tools at varying video sampling rates (i.e. 1, 5, 25 frames per seconds FPS). A demonstration is included in the qualitative video.

All frames	HOTA	DetA	DetPr	DetRe	LocA	AssA	AssPr	AssRe
	60.2	71.7	83.1	77.9	86.5	50.6	71.7	55.4

<p>(a) None</p>	<p>HOTA 67.2</p> <p>DetA 71.5</p> <p>DetPr 83.0</p> <p>DetRe 81.4</p> <p>LocA 87.3</p> <p>AssA 63.4</p> <p>AssPr 76.7</p> <p>AssRe 71.1</p>	<p>(b) Bleeding from injuries</p>	<p>HOTA 63.2</p> <p>DetA 72.9</p> <p>DetPr 83.5</p> <p>DetRe 78.5</p> <p>LocA 86.5</p> <p>AssA 55.1</p> <p>AssPr 76.7</p> <p>AssRe 59.6</p>	<p>(c) Crowded scene</p>	<p>HOTA 57.0</p> <p>DetA 70.2</p> <p>DetPr 83.6</p> <p>DetRe 75.3</p> <p>LocA 86.1</p> <p>AssA 46.5</p> <p>AssPr 72.4</p> <p>AssRe 51.0</p>
<p>(d) Smoke from coagulation</p>	<p>HOTA 54.3</p> <p>DetA 61.8</p> <p>DetPr 82.0</p> <p>DetRe 66.8</p> <p>LocA 86.0</p> <p>AssA 48.0</p> <p>AssPr 70.3</p> <p>AssRe 52.1</p>	<p>(e) Occlusion</p>	<p>HOTA 58.5</p> <p>DetA 71.5</p> <p>DetPr 83.4</p> <p>DetRe 77.2</p> <p>LocA 86.4</p> <p>AssA 48.0</p> <p>AssPr 70.8</p> <p>AssRe 52.6</p>	<p>(f) Fast motion blurring</p>	<p>HOTA 50.4</p> <p>DetA 55.3</p> <p>DetPr 75.0</p> <p>DetRe 61.9</p> <p>LocA 82.9</p> <p>AssA 46.6</p> <p>AssPr 72.0</p> <p>AssRe 53.2</p>
<p>(g) Camera Lens Fouling</p>	<p>HOTA 54.2</p> <p>DetA 65.7</p> <p>DetPr 78.7</p> <p>DetRe 71.6</p> <p>LocA 83.3</p> <p>AssA 49.6</p> <p>AssPr 71.1</p> <p>AssRe 61.0</p>	<p>(h) Under coverage, trocar view</p>	<p>HOTA 57.2</p> <p>DetA 68.3</p> <p>DetPr 81.1</p> <p>DetRe 76.6</p> <p>LocA 87.5</p> <p>AssA 48.0</p> <p>AssPr 69.1</p> <p>AssRe 55.9</p>	<p>(i) Specular light reflection</p>	<p>HOTA 46.7</p> <p>DetA 54.9</p> <p>DetPr 72.8</p> <p>DetRe 64.0</p> <p>LocA 84.0</p> <p>AssA 40.7</p> <p>AssPr 65.2</p> <p>AssRe 54.1</p>

**Fig. 5.** Performance assessment of SurgiTrack amidst surgical visual challenges. Overall performance is tabulated at the top, preceded by quantitative and qualitative results showcasing tracking performance on specific visual challenge frames. Values in black denote comparable performance (within the average range,  $\pm 1.0$ ). Values in green indicate above-average performance, while red values indicate decreasing performance below average. The breakdown explores distinct tracking metrics focusing on detection, localization, and association or re-identification. A demo is included in the qualitative results video.

### 6.6. Analysis of multi-class tracking results

We analyze in Table 7 the tracking performance per tool class and observe that hook is the best tracked with a HOTA above 75.6% for the proposed solutions. This class-agnostic tracking result reveals that regularly used tools such as bipolar and clipper has a medium tracking accuracy. Rarely used tools (e.g. scissors, irrigator) have the least scores. The grasper with the most multiple instances is very difficult to track, similarly the specimen bag is badly tracked due to its constantly deforming shapes.

Revealed by this analysis, is the impact of the infusing surgical knowledge in the existing trackers, via the KB baselines. This greatly improves the tracking of multiple instances tools like grasper and marginal improves on tracking of single instance tools. Meanwhile, direction features learning the infusion of tool category and direction eliminates the need for hard-coded domain knowledge, giving a better performance. With superior performance over appearance-based features in differentiating tools with identical appearances but different directions (e.g., graspers), we demonstrate that the learned features effectively capture direction-specific semantics, validating their designation as “direction” features.

### 6.7. Impact of direction estimation on tracking FPS

We measure the impact of direction estimation on variable video sampling rates and show that it facilitates track re-identification in the events of non-overlapping bounding boxes as shown in Fig. 4. Unlike the compared models which are unable to track the tools at lower video sampling rates, SurgiTrack maintains a relatively stable tracking performance sampled across 1, 5, and 25 FPS in all the trajectory perspectives. This result justifies the re-identification strength of the proposed direction features and its modeling as a proxy to the tool operators, initially identified as the most promising cue for differentiating tool instances of the same class.

### 6.8. Tracking evaluation under surgical visual challenges

We analyze the tool tracking model across various surgical conditions, measured through HOTA metrics in Fig. 5, to uncover insights into how the model interacts with the complexities of the surgical environment. The overall HOTA score across all frames for SurgiTrack (SSL) peaks at 60%, providing a benchmark for comparison. Each visual condition is scrutinized based on its impact on different aspects of



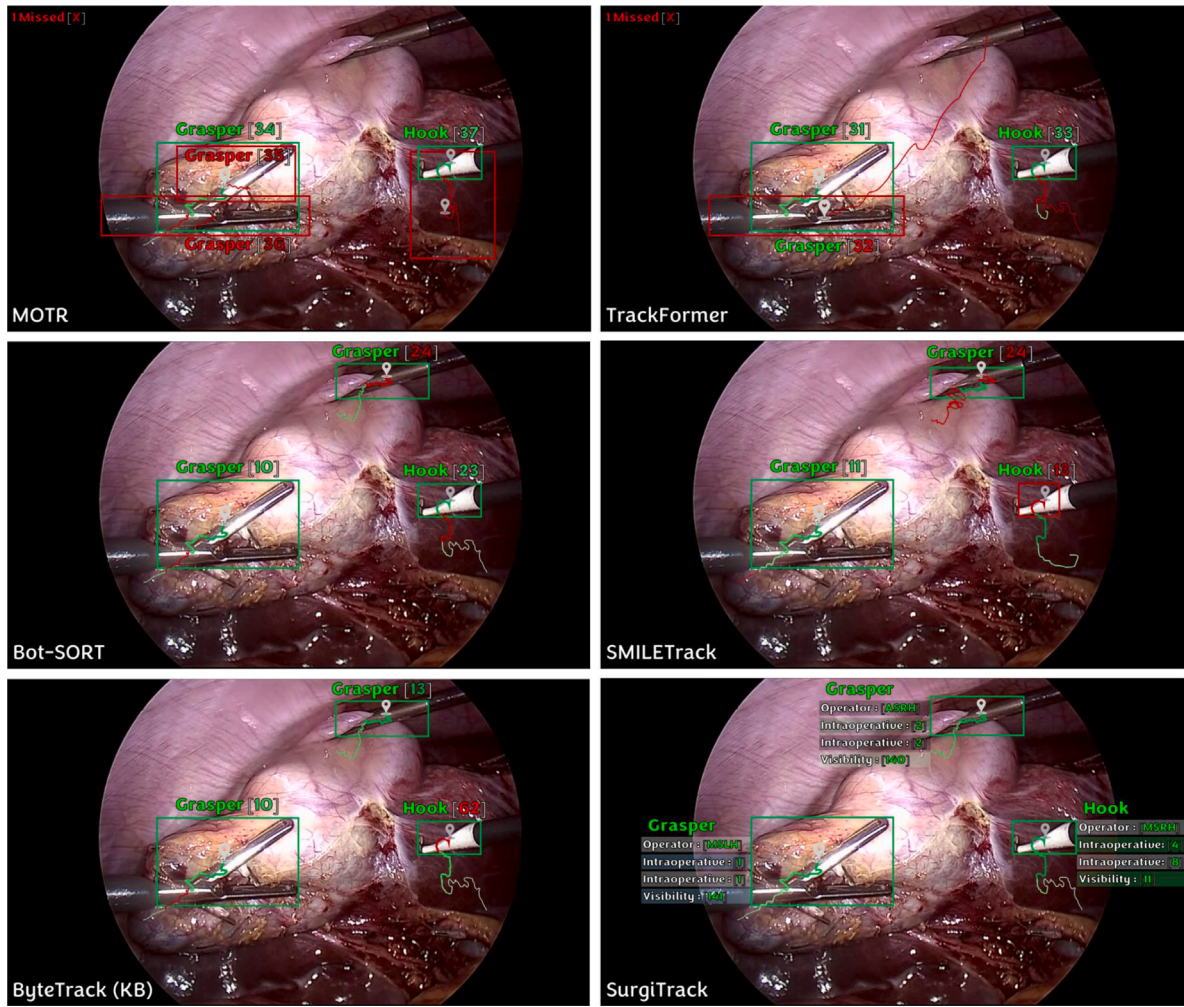


Fig. 6. Qualitative result of SurgiTrack in comparison with some existing methods. Bounding box represents tool detection, tool name represents tool classification, number in block parenthesis represents track identity, and scribble represents tracklet (max. 2 s). Green color indicates correctness, red indicates failure. A demo is included in the qualitative results video.

tracking, encompassing detection (accuracy, precision, recall), localization accuracy, and re-identification/association (accuracy, precision, recall).

Under clear scenes, devoid of visual challenges, the model excels with a substantial increase in performance (+7.0% HOTA, +12.8% AssA), showcasing efficient tool discernment, precise localization, and seamless tracking. In scenarios with bleeding, the model demonstrates resilience, maintaining stable detection and localization, and showing improved association. This suggests the model's capability to detect and locate tools even amidst dynamic visual changes and potential tool discoloration from blood stains, which may have been treated as a form of augmentation owing to the abundant bleeding case samples.

Conversely, in crowded scenes, the model faces challenges as numerous tools closely packed together hinder detection accuracy, potentially affecting accurate tracking. Disentangling cluttered tools proves difficult, potentially affecting accurate identity matching (−4.1% AssA, −4.4% AssRe). Smoke from coagulation tools decreases the tracking score due to impaired visibility, resulting in significant missed detections (−9.9% DetA, −11.1% DetRe) and diminishing association accuracy.

The detection and localization accuracies remains unaffected in the cases of occlusion likely due to groundtruth bounding boxes including only the unoccluded tool regions. However, the occluded areas, which are often the discriminative tooltips, pose challenges for tool classification and re-identification tasks (−2.6% AssA, −2.8% AssRe), slightly

impacting tracking performance. Rapid motion-induced blurring introduces uncertainties in tool appearances, restricting the model's ability to accurately identify and track instruments, significantly reducing the HOTA score by −9.8% and impacting all tracking aspects significantly. Fluids such as blood and bile sometimes fouds the camera lens, impair visibility, and disrupt consistent identity matching. This negatively impact tracking in all aspects with a 6.0% reduction in HOTA score.

Limited camera coverage poses a major challenge, causing numerous missed detections, although localizing focused tools remains precise. Constrained trocar view hampers the model's ability to track and manage tool identities, especially complicating the interpretation of tool direction, as a link to the operating surgeon's hand. Specular light reflection emerges as the most impactful visual challenge, leading to a notable decrease in all tracking aspects, with a huge loss of 13.5% points in HOTA. The excessive brightness from specular light obscures visual details, leading to a 16.8% and 9.9% reduction in detection and association accuracies respectively. Though rare in this dataset, appearance changes in a single tool can also result from the use of electrocautery instruments.

This analysis sheds light on how visual challenges during surgery impact vision tasks such as tool tracking. While some pose significant hurdles, others inadvertently offer clear visual cues that help track tools. The insights can inform future enhancements to improve tool tracking in challenging surgical environment.

Beyond the failure modes due to the surgical visual challenges, we observe that the model experiences difficulties at the beginning of a






Fig. 7. Qualitative result of SurgiTrack in comparison with a state of the art method (BotSORT) on tracking across variable frame rates (1FPS, 5FPS and 25FPS). Tick blue bounding box represents tool detection at current time. dotted gray bounding boxes detection at previous times, tool name represents tool classification, track identity number is written above each box. A demo is included in the qualitative results video.

video due to limited historical data for the HGBM to rely on, leading to more frequent identity fragmentations, however, this usually stabilizes in few seconds.

As a limitation, our approach would not detect a replacement of a tool by another tool of the same class (and same manufacturer) if operated via the same trocar port. This is because the proposed direction feature is designed to uniquely differentiate tools based on their originating directions — the trocar ports. Additional signals would be needed to tackle this situation. The use of direction features may face challenges in certain type of surgery such as endoscopic endonasal surgery where the nostrils are closely space, or where trocars are closely spaced with little or no directional difference.

#### 6.9. Qualitative results and demos

In Fig. 6, we present a qualitative analysis of our method alongside some of the existing approaches, complementing our quantitative

evaluation. Performance-wise, transformer-based solutions (e.g., MOTR, TrackFormer) struggle with precise tool localization. Location-based methods (e.g., ByteTrack) mostly encounter identity mismatch issues. Appearance-based trackers (e.g., Bot-SORT) demonstrate strengths but sometimes misclassify tools and miss detections, akin to Similarity-based trackers such as SMILETrack. Our method, SurgiTrack, displays enhanced resilience in tracking tools jointly across multiple perspectives compared to existing methods. Fig. 7 demonstrates SurgiTrack's robustness to variable frame rates (1, 5, 25 fps). This highlights the impact of direction re-ID features compared to the appearance features in the state-of-the-art method (BotSORT), which frequently assigns new track IDs when there is insufficient overlap between bounding boxes.

**Video Demonstration:** Given that tracking is better appreciated in motion, as an appendix, we provide a full video demonstrating in detail the qualitative performance of our method across multiple evaluation settings. The demo video is available at <https://vimeo.com/951853260>, and as supplementary material.

## 7. Conclusion

In this work, we propose, SurgiTrack, a novel deep learning approach for multi-class multi-tool tracking in surgical videos. Our approach utilizes an attention-based deep learning model for tool identity association by learning the tool motion direction which we conceived as a proxy to linking the tools to the operating surgeons' hands via the trocars. We demonstrate that the motion direction features are superior to location, appearance, and similarity features for the re-identification of surgical tools given the non-distinctiveness of most tools' appearance, especially the ones from the same or similar classes. We show that the direction features can be learnt in 3 different paradigms of full-, weak-, and self-supervision depending on the availability of training labels. We also design a harmonizing bipartite matching graph to enable non-conflicting and synchronized tracking of tools across three perspectives of intraoperative, intracorporeal, and visibility within the camera field of view, which represent the various ways of considering the temporal duration of a tool trajectory. Additionally, we benchmark several deep learning methods for tool detection and tracking on the newly introduced CholecTrack20 dataset and conducted ablation studies on the suitability of existing re-identification features for accurate tool tracking. Our proposed model emerges as a promising solution for multi-class multi-tool tracking in surgical procedures, showcasing adaptability across different training paradigms and demonstrating strong performance in essential tracking metrics. We also evaluate our model across different surgical visual challenges such as bleeding, smoke, occlusion, camera fouling, light reflection, etc., and presents insightful findings on their impact on visual tracking in surgical videos. Qualitative results also show that our method is effective in handling challenging situations compare to the baselines and can effortlessly track tools irrespective of the video frame sampling rate.

## CRediT authorship contribution statement

**Chinedu Innocent Nwoye:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Nicolas Padoy:** Writing – review & editing, Supervision, Project administration, Funding acquisition.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported by French state funds managed within the Plan Investissements d'Avenir by the ANR under references: National AI Chair AI4ORSafety [ANR-20-CHIA-0029-01], IHU Strasbourg [ANR-10-IAHU-02] and by BPI France [Project 5G-OR]. This work was granted access to the servers/HPC resources managed by CAMMA, IHU Strasbourg, Unistra Mesocentre, and GENCI-IDRIS [Grant 2021-AD011011638R3, 2021-AD011011638R4]. The authors acknowledge the suggestions of Rupak Bose on the visualization plots and thank Dr. Joël L. Lavanchy for some useful clinical discussions.

## Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.media.2024.103438>.

## Data availability

Data will be made available on request.

## References

- Aharon, N., Orfaig, R., Bobrovsky, B.Z., 2022. Bot-sort: Robust associations multi-pedestrian tracking. *arXiv preprint arXiv:2206.14651*.
- Alsheakhali, M., Yigitsoy, M., Eslami, A., Navab, N., 2015. Surgical tool detection and tracking in retinal microsurgery. In: *Medical Imaging 2015: Image-Guided Procedures, Robotic Interventions, and Modeling*. SPIE, pp. 245–250.
- Bahdanau, D., 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Behrens, A., Grimm, J., Gross, S., Aach, T., 2011. Inertial navigation system for bladder endoscopy. In: *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. pp. 5376–5379. <http://dx.doi.org/10.1109/IEMBS.2011.6091330>.
- Bergmann, P., Meinhardt, T., Leal-Taixe, L., 2019. Tracking without bells and whistles. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 941–951.
- Bernardin, K., Stiefel, R., 2008. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP J. Image Video Process.* 2008, 1–10.
- Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B., 2016. Simple online and realtime tracking. In: *2016 IEEE International Conference on Image Processing. ICIP, IEEE*, pp. 3464–3468.
- Bouget, D., Allan, M., Stoyanov, D., Jannin, P., 2017. Vision-based and marker-less surgical tool detection and tracking: a review of the literature. *Med. Image Anal.* 35, 633–654.
- Cartucho, J., Wang, C., Huang, B., Elson, D.S., Darzi, A., Giannarou, S., 2022. An enhanced marker pattern that achieves improved accuracy in surgical tool tracking. *Comput. Methods Biomech. Biomed. Eng.: Imaging Vis.* 10, 400–408.
- Chen, C.F.R., Fan, Q., Panda, R., 2021. Crossvit: Cross-attention multi-scale vision transformer for image classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 357–366.
- Chmarr, M.K., Grimbergen, C., Dankelman, J., 2007. Systems for tracking minimally invasive surgical instruments. *Minim. Invasive Ther. Allied Technol.* 16, 328–340.
- Chu, P., Wang, J., You, Q., Ling, H., Liu, Z., 2023. Transmot: Spatial-temporal graph transformer for multiple object tracking. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 4870–4880.
- Danelljan, M., Bhat, G., Shahbaz Khan, F., Felsberg, M., 2017. Eco: Efficient convolution operators for tracking. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6638–6646.
- Dendorfer, P., Rezatofighi, H., Milan, A., Shi, J., Cremers, D., Reid, I., Roth, S., Schindler, K., Leal-Taixé, L., 2020. Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv preprint arXiv:2003.09003*.
- Dockter, R., Sweet, R., Kowalewski, T., 2014. A fast, low-cost, computer vision approach for tracking surgical tools. In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE*, pp. 1984–1989.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al., 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Du, X., Allan, M., Dore, A., Ourselin, S., Hawkes, D., Kelly, J.D., Stoyanov, D., 2016. Combined 2d and 3d tracking of surgical instruments for minimally invasive and robotic-assisted surgery. *Int. J. Comput. Assist. Radiol. Surg.* 11, 1109–1119.
- Du, Y., Wan, J., Zhao, Y., Zhang, B., Tong, Z., Dong, J., 2021. Gaiotracker: A comprehensive framework for mcmot with global information and optimizing strategies in visdrone 2021. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2809–2819.
- Du, Y., Zhao, Z., Song, Y., Zhao, Y., Su, F., Gong, T., Meng, H., 2023. Strongsort: Make deepsort great again. *IEEE Trans. Multimed.*
- Dutkiewicz, P., Kietczewski, M., Kowalski, M., Wroblewski, W., 2005. Experimental verification of visual tracking of surgical tools. In: *Proceedings of the Fifth International Workshop on Robot Motion and Control, 2005. RoMoCo'05, IEEE*, pp. 237–242.
- Fathollahi, M., Sarhan, M.H., Pena, R., DiMonte, L., Gupta, A., Ataliwala, A., Barker, J., 2022. Video-based surgical skills assessment using long term tool tracking. In: *Medical Image Computing and Computer Assisted Intervention-MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VII*. Springer, pp. 541–550.
- Fried, M.P., Kleefield, J., Gopal, H., Reardon, E., Ho, B.T., Kuhn, F.A., 1997. Image-guided endoscopic surgery: results of accuracy and performance in a multi-center clinical study using an electromagnetic tracking system. *Laryngoscope* 107, 594–601.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 770–778.
- Huang, B., Tsai, Y.Y., Cartucho, J., Vyas, K., Tuch, D., Giannarou, S., Elson, D.S., 2020. Tracking and visualization of the sensing area for a tethered laparoscopic gamma probe. *Int. J. Comput. Assist. Radiol. Surg.* 15, 1389–1397.
- Jo, K., Im, J., Kim, J., Kim, D.S., 2017. A real-time multi-class multi-object tracker using yolov2. In: *2017 IEEE International Conference on Signal and Image Processing Applications. ICSIPA, IEEE*, pp. 507–511.
- Kalman, R.E., 1960. A new approach to linear filtering and prediction problems.



- Kuhn, H.W., 1955. The hungarian method for the assignment problem. *Nav. Res. Logist. Q.* 2, 83–97.
- Lee, B., Erdenee, E., Jin, S., Nam, M.Y., Jung, Y.G., Rhee, P.K., 2016. Multi-class multi-object tracking using changing point detection. In: *Computer Vision–ECCV 2016 Workshops: Amsterdam, the Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part II 14*. Springer, pp. 68–83.
- Lee, E.J., Plishker, W., Liu, X., Bhattacharyya, S.S., Shekhar, R., 2019. Weakly supervised segmentation for real-time surgical tool tracking. *Healthc. Technol. Lett.* 6, 231–236.
- Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixé, L., Leibe, B., 2021. Hota: A higher order metric for evaluating multi-object tracking. *Int. J. Comput. Vis.* 129, 548–578.
- Ma, L., Wang, J., Kiyomatsu, H., Tsukihara, H., Sakuma, I., Kobayashi, E., 2021. Surgical navigation system for laparoscopic lateral pelvic lymph node dissection in rectal cancer surgery using laparoscopic-vision-tracked ultrasonic imaging. *Surg. Endosc.* 35, 6556–6567.
- Maggiolino, G., Ahmad, A., Cao, J., Kitani, K., 2023. Deep oc-sort: Multi-pedestrian tracking by adaptive re-identification. *arXiv preprint arXiv:2302.11813*.
- Martin-Gomez, A., Li, H., Song, T., Yang, S., Wang, G., Ding, H., Navab, N., Zhao, Z., Armand, M., 2023. Sttar: surgical tool tracking using off-the-shelf augmented reality head-mounted displays. *IEEE Trans. Vis. Comput. Graphics*.
- Meinhardt, T., Kirillov, A., Leal-Taixe, L., Feichtenhofer, C., 2022. Trackformer: Multi-object tracking with transformers. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8844–8854.
- Nwoye, C.I., 2021. Deep Learning Methods for the Detection and Recognition of Surgical Tools and Activities in Laparoscopic Videos (Ph.D. thesis). Université de Strasbourg.
- Nwoye, C.I., Elgohary, K., Srinivas, A., Zaid, F., Lavanchy, J.L., Padoy, N., 2023. CholecTrack20: A dataset for multi-class multiple tool tracking in laparoscopic surgery. *arXiv preprint*.
- Nwoye, C.I., Elgohary, K., Srinivas, A., Zaid, F., Lavanchy, J.L., Padoy, N., 2024. CholecTrack20 dataset for multi-class multiple tool tracking in laparoscopic surgery. *Synapse* <http://dx.doi.org/10.7303/syn53182642>.
- Nwoye, C.I., Mutter, D., Marescaux, J., Padoy, N., 2019. Weakly supervised convolutional lstm approach for tool tracking in laparoscopic videos. *Int. J. Comput. Assist. Radiol. Surg.* 14, 1059–1067.
- Pedrett, R., Mascagni, P., Beldi, G., Padoy, N., Lavanchy, J.L., 2023. Technical skill assessment in minimally invasive surgery using artificial intelligence: a systematic review. *Surg. Endosc.* 1–13.
- Pezementi, Z., Voros, S., Hager, G.D., 2009. Articulated object tracking by rendering consistent appearance parts. In: *2009 IEEE International Conference on Robotics and Automation*. IEEE, pp. 3940–3947.
- Reiter, A., Allen, P.K., Zhao, T., 2012. Feature classification for tracking articulated surgical tools. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012: 15th International Conference, Nice, France, October 1–5, 2012, Proceedings, Part II 15*. Springer, pp. 592–600.
- Reiter, A., Allen, P.K., Zhao, T., 2014. Appearance learning for 3d tracking of robotic surgical tools. *Int. J. Robot. Res.* 33, 342–356.
- Richa, R., Balicki, M., Meisner, E., Sznitman, R., Taylor, R., Hager, G., 2011. Visual tracking of surgical tools for proximity detection in retinal surgery. In: *Information Processing in Computer-Assisted Interventions: Second International Conference, IPCAI 2011, Berlin, Germany, June 22 2011. Proceedings 2*. Springer, pp. 55–66.
- Rieke, N., Tan, D.J., Alsheikhali, M., Tombari, F., di San Filippo, C.A., Belagiannis, V., Eslami, A., Navab, N., 2015. Surgical tool tracking and pose estimation in retinal microsurgery. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part I 18*. Springer, pp. 266–273.
- Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C., 2016. Performance measures and a data set for multi-target, multi-camera tracking. In: *European Conference on Computer Vision*. Springer, pp. 17–35.
- Robu, M., Kadkhodamohammadi, A., Luengo, I., Stoyanov, D., 2021. Towards real-time multiple surgical tool tracking. *Comput. Methods Biomech. Biomed. Eng.: Imaging Vis.* 9, 279–285.
- Rueckert, T., Rueckert, D., Palm, C., 2023. Methods and datasets for segmentation of minimally invasive surgical instruments in endoscopic images and videos: A review of the state of the art. *arXiv preprint arXiv:2304.13014*.
- Shao, S., Zhao, Z., Li, B., Xiao, T., Yu, G., Zhang, X., Sun, J., 2018. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*.
- Song, H., Yang, S., Wu, Z., Moradi, H., Taylor, R.H., Kang, J.U., Salcudean, S.E., Bocktor, E.M., 2023. Arc-to-line frame registration method for ultrasound and photoacoustic image-guided intraoperative robot-assisted laparoscopic prostatectomy. *arXiv preprint arXiv:2306.12590*.
- Speidel, S., Sudra, G., Senemaud, J., Dretnschew, M., Müller-Stich, B.P., Gutt, C., Dillmann, R., 2008. Recognition of risk situations based on endoscopic instrument tracking and knowledge based situation modeling. In: *Medical Imaging 2008: Visualization, Image-Guided Procedures, and Modeling*. SPIE, pp. 326–333.
- Sun, P., Cao, J., Jiang, Y., Zhang, R., Xie, E., Yuan, Z., Wang, C., Luo, P., 2020. Transtrack: Multiple object tracking with transformer. *arXiv preprint arXiv:2012.15460*.
- Sznitman, R., Ali, K., Richa, R., Taylor, R.H., Hager, G.D., Fua, P., 2012. Data-driven visual tracking in retinal microsurgery. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012: 15th International Conference, Nice, France, October 1–5, 2012, Proceedings, Part II 15*. Springer, pp. 568–575.
- Tan, M., Le, Q., 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International Conference on Machine Learning*. PMLR, pp. 6105–6114.
- Wang, C.Y., Bochkovskiy, A., Liao, H.Y.M., 2023a. Yolo7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7464–7475.
- Wang, Y.H., Hsieh, J.W., Chen, P.Y., Chang, M.C., So, H.H., Li, X., 2023b. Smiletrack: Similarity learning for occlusion-aware multiple object tracking. *arXiv:2211.08824*.
- Wang, Y., Sun, Q., Liu, Z., Gu, L., 2022. Visual detection and tracking algorithms for minimally invasive surgical instruments: A comprehensive review of the state-of-the-art. *Robot. Auton. Syst.* 149, 103945.
- Wang, Z., Zheng, L., Liu, Y., Li, Y., Wang, S., 2020. Towards real-time multi-object tracking. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*. Springer, pp. 107–122.
- Wojke, N., Bewley, A., Paulus, D., 2017. Simple online and realtime tracking with a deep association metric. In: *2017 IEEE International Conference on Image Processing*. ICIP, IEEE, pp. 3645–3649.
- Xu, H., Runciman, M., Cartucho, J., Xu, C., Giannarou, S., 2023. Graph-based pose estimation of texture-less surgical tools for autonomous robot control. In: *2023 IEEE International Conference on Robotics and Automation*. ICRA, IEEE, pp. 2731–2737.
- Xu, L., Zhang, H., Wang, J., Li, A., Song, S., Ren, H., Qi, L., Gu, J.J., Meng, M.Q.H., 2022. Information loss challenges in surgical navigation systems: From information fusion to ai-based approaches. *Inf. Fusion*.
- Ye, M., Zhang, L., Giannarou, S., Yang, G.Z., 2016. Real-time 3d tracking of articulated tools for robotic surgery. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part I 19*. Springer, pp. 386–394.
- Zeng, F., Dong, B., Zhang, Y., Wang, T., Zhang, X., Wei, Y., 2022. Motr: End-to-end multiple-object tracking with transformer. In: *European Conference on Computer Vision*. Springer, pp. 659–675.
- Zhang, Y., Sun, P., Jiang, Y., Yu, D., Weng, F., Yuan, Z., Luo, P., Liu, W., Wang, X., 2022. Bytetrack: Multi-object tracking by associating every detection box. In: *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*. Springer, pp. 1–21.
- Zhang, Y., Wang, C., Wang, X., Zeng, W., Liu, W., 2021. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *Int. J. Comput. Vis.* 129, 3069–3087.
- Zhao, Z., Chen, Z., Voros, S., Cheng, X., 2019a. Real-time tracking of surgical instruments based on spatio-temporal context and deep learning. *Comput. Assist. Surg.* 24, 20–29.
- Zhao, Z., Voros, S., Chen, Z., Cheng, X., 2019b. Surgical tool tracking based on two cnns: from coarse to fine. *J. Eng.* 2019, 467–472.
- Zhu, P., Wen, L., Du, D., Bian, X., Fan, H., Hu, Q., Ling, H., 2021. Detection and tracking meet drones challenge. *IEEE Trans. Pattern Anal. Mach. Intell.* 44, 7380–7399.