

A Kinematic Bottleneck Approach For Pose Regression of Flexible Surgical Instruments directly from Images

Luca Sestini^{1,2}, Benoit Rosa¹, Elena De Momi², Giancarlo Ferrigno² and Nicolas Padoy^{1,3}

Abstract—3-D pose estimation of instruments is a crucial step towards automatic scene understanding in robotic minimally invasive surgery. Although robotic systems can potentially directly provide joint values, this information is not commonly exploited inside the operating room, due to its possible unreliability, limited access and the time-consuming calibration required, especially for continuum robots. For this reason, standard approaches for 3-D pose estimation involve the use of external tracking systems. Recently, image-based methods have emerged as promising, non-invasive alternatives. While many image-based approaches in the literature have shown accurate results, they generally require either a complex iterative optimization for each processed image, making them unsuitable for real-time applications, or a large number of manually-annotated images for efficient learning. In this paper we propose a self-supervised image-based method, exploiting, at training time only, the imprecise kinematic information provided by the robot. In order to avoid introducing time-consuming manual annotations, the problem is formulated as an auto-encoder, smartly *bottlenecked* by the presence of a physical model of the robotic instruments and surgical camera, forcing a separation between image background and kinematic content. Validation of the method was performed on *semi-synthetic*, *phantom* and *in-vivo* datasets, obtained using a flexible robotized endoscope, showing promising results for real-time image-based 3-D pose estimation of surgical instruments.

Index Terms—Computer Vision for Medical Robotics; Surgical Robotics; Laparoscopy; Deep Learning Methods; Self-Supervision

I. INTRODUCTION

ROBOTIC Minimally Invasive Surgery (RMIS) is becoming a standard paradigm across different surgical specialties [1], providing surgeons with a significantly less invasive method of accessing distal sites, together with a level

Manuscript received: October 15, 2020; Revised: January 10, 2021; Accepted: February 8, 2021.

This paper was recommended for publication by Editor Eric Marchand upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported by the ATLAS project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 813782. This work is also partially supported by French State Funds managed by the Agence Nationale de la Recherche (ANR) through the Investissements d'Avenir Program under Grant ANR-11-LABX-0004 (Labex CAMI) and Grant ANR-10-IAHU-02 (IHU-Strasbourg)

¹ Luca Sestini, Benoit Rosa and Nicolas Padoy are with ICube, University of Strasbourg, CNRS, IHU Strasbourg, France {sestini,b.rosa,npadoy}@unistra.fr

² Luca Sestini, Elena De Momi and Giancarlo Ferrigno are with Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milano, Italy {elena.demomi,giancarlo.ferrigno}@polimi.it

³ Nicolas Padoy is with IHU Strasbourg, Strasbourg, France.

Digital Object Identifier (DOI): see top of this page.

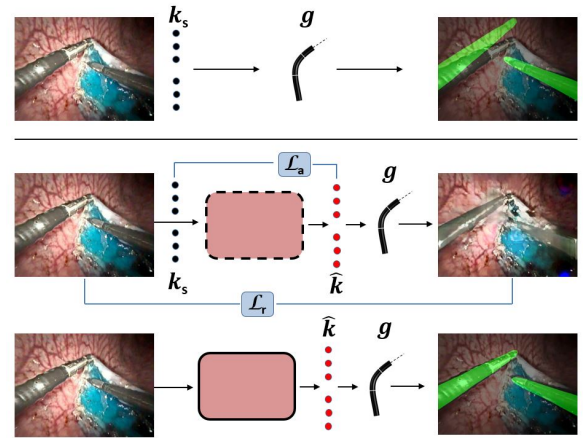


Fig. 1. Illustration of the proposed approach. Top row: example of an image with associated kinematics k_s , measured by the robotic system. Using a model of the robotic instruments, k_s can be mapped to their 3-D shape g , and then reprojected on the image plane. The projected silhouette shows the inaccuracy of k_s , due to the strong tool-tissue interaction; middle row: our model (red box) is trained by means of a self-supervised reconstruction loss (\mathcal{L}_r) and a weakly-supervised loss (\mathcal{L}_a) provided by the uncertain kinematics k_s , without any manual supervision; bottom row: at inference time the model predicts kinematics from image-information only, with a superior accuracy with respect to the measured kinematics k_s .

of precision and dexterity comparable to open surgery. Robotic systems are generally integrated with computer assistance, providing the surgeons with rich information about the intervention. A necessary step towards complete integration of computer assistance and RMIS is represented by the automatic localization of tools in the surgical scene, both at image level and in the 3-D space [2]. Tool localization at image level offers a wide range of applications, from workflow analysis to tool subtraction in panorama reconstruction, augmented reality and visual servoing [3]. At 3-D level, localization of the instruments can be exploited to provide dynamic motion constraints, especially in the context of flexible robots [4][5], in haptic-feedback applications [6], in tumor localization via remote palpation [7], to analyze surgeons skills and in general for the above-mentioned 2-D applications, via reprojection onto the image space. In this paper we concentrate on the 3-D pose estimation problem, focusing, but not limiting to, continuum robots.

Continuum robots have been developed to access the most challenging surgical sites in a minimally invasive way [8]. Compared to classical rigid articulated robots, continuum




Fig. 2. Full training architecture: given an input image x and the corresponding imprecise measured kinematics k_s , the training architecture forces a separation between the appearance of the image and its kinematic content, while keeping a self-supervised formulation. In the bottom branch, the module β transforms x into a *backgroundized* version of itself x_b . In the top branch, the regressor ψ reduces x to the estimated kinematic configuration \hat{k} , which is then converted by the physical model ρ of the instruments and surgical camera in the binary silhouette projection \hat{m} of the instruments on the image plane. A decoder ϕ tries to reconstruct the input image x from x_b and \hat{m} . The model is trained by means of the reconstruction loss \mathcal{L}_r , helped by the auxiliary loss \mathcal{L}_a . The *backgroundizer* and the *physical module* are trained in advance and frozen during the training of regressor and decoder.

robots bend continuously when actuated, resulting in increased dexterity, at the price of more challenging control and modelling [9]. Recently developed robotic systems can potentially provide the kinematic information associated with each joint, measuring it at the motor side. However, this information is often unreliable due to two main factors: 1) tool-tissue interactions, which can modify the shape of a flexible instruments; 2) un-modelled non-linearities (e.g. cable friction, slack in instrument channel, backlash), which can lead to loss of motion between the motors and the instruments; these two factors can lead to a non-deterministic mapping between measured kinematics and actual robot configuration [10] (fig. 1, top row). Some of the mentioned phenomena depend on the type of instruments used, as well as their initial positioning, usually performed manually, making calibration a necessary, time-consuming step before every procedure. For those reasons accessing robot kinematics in real-time inside the operating room (OR) is uncommon [11], even when access to the robot API is granted. Standard methods for robots shape estimation involve attaching external sensors to the instruments (mainly electromagnetic, optical markers, Fiber Bragg Gratings), using trackers to estimate the shape [9][12][13][14]. However, relying on external sensors involves several undesirable actions to be taken, such as modifying the instrument design to fit the sensors, introducing trackers in the OR and in general, altering the normal surgical workflow. For this reason, vision-based marker-less solutions have recently gained strong interest [15].

The problem of estimating robot instruments pose from surgical camera images is generally posed as maximizing the agreement between a parametrized model of the instruments and features extracted from the image. Statistical region-based methods formulate the problem as estimating the pose parameters such that the corresponding silhouette regions, determined by the projection of a 3-D model of the instruments on the camera plane, match the silhouette regions estimated from a pixel-wise classification of the image. Works on rigid articulated robots, based on this approach, have shown good

accuracy [16][11]. However, those methods require a per-frame optimization, which limits their real-time applicability. In addition, their accuracy is dependent on the quality of the segmentation model: the proposed Random Forest algorithms, initialized from a single frame, may not be able to handle a high variability of the data appearance (although online learning is implemented in [11]), and more powerful models may be needed, requiring, in turn, a large amount of annotated data for training. A different approach for surgical instruments pose estimation has been proposed in [17][18][19], which used deep learning models to directly localize the joint position on the image. Despite achieving high localization accuracy, those methods are limited to 2-D space, and require manual annotations for training. In general, for successful generalizability, supervised machine learning approaches require a great amount of labelled data acquired through manual annotation, which is both costly and time-consuming.

Several works in the computer-assisted interventions community have recently explored ways to mitigate the need for such amount of manual annotations, mainly focusing on two approaches: developing strategies to learn from synthetic data [20][21][22] and reducing the amount of labels needed for training [23][24]. In [25], Nwoye et al. explored weak-supervision, using cheap frame-level annotations on tool presence, to train a model for 2-D tool tracking. In [26], Tanwani et al. developed a semi-supervised approach based on surgical actions grouping, in order to obtain meaningful embedded representations of endoscopic video segments, showing promising results for downstream tasks as segmentation and pose imitation. A side of weak-supervision that has been marginally explored, especially in the surgical domain, is represented by the use of easy-to-obtain, but possibly inaccurate information. This paper demonstrates the benefits of using inaccurate kinematics information for self-supervision during training.

In this work we address the problem of estimating the 3-D pose of flexible surgical instruments and propose a new self-supervised, image-based approach, which:

- can estimate the *full* 3-D pose of instruments, parametrized by means of their joint values, from images only, with no additional optimization required at inference time, making it easily suitable for real-time applications;
- is trained end-to-end in a self-supervised way, exploiting an auto-encoder formulation, smartly *bottlenecked* by the presence of a physical model of instruments and surgical camera, in order to extract their kinematics;
- leverages, at training time only, the rough localization provided by the uncertain measured kinematics, without requiring any manual annotation.

II. METHOD

Our goal is to train a regressor model to estimate the pose of surgical instruments from image data acquired by the surgical camera. Differently from standard methods as [11], where the estimated pose is represented by the position and orientation of the end-effector, the regressor is trained to directly estimate the value of the n kinematic joints $\mathbf{k} = \{k_0, k_1, \dots, k_{n-1}\}$ of the instruments. The joint values, combined with a geometric forward kinematic model of the instruments, allow to reconstruct their *full* 3-D shape.

The problem is formulated as learning the set of parameters θ_ψ of the regressor model $\psi : x \rightarrow \hat{\mathbf{k}}$ that maps an image x , containing the instruments in a configuration represented by the set of ground truth joint values $\mathbf{k} = \{k_0, k_1, \dots, k_{n-1}\}$, to the set of estimated joint values $\hat{\mathbf{k}}$, such that:

$$\hat{\theta}_\psi = \arg \min_{\theta_\psi} (\psi(\theta_\psi, x) - \mathbf{k}). \quad (1)$$

However, \mathbf{k} is not known in practice. The measured joint values \mathbf{k}_s , recorded by the robotic system, are generally inaccurate, due to tool-tissue interactions and possible un-modelled non-linearities, thus they cannot be directly used as supervision signal. In order to avoid relying on time-consuming and expensive manual annotations, we propose a training architecture built as an auto-encoder, trained by reconstructing the input image x as output of the decoder model ϕ , defined by the set of parameters θ_ϕ . The complete training architecture can be observed in fig. 2. The architecture, inspired by [27], is based on the idea of separating the appearance of the image from its kinematic content. In order to do that, the architecture presents two separate branches. The *appearance* branch consists of a module β , called *backgroundizer module*, which converts the input image x to an approximate background-only version of itself x_b . The *content* branch contains the kinematic regressor ψ . However, without further constraints, the kinematic extraction would not be possible: the regressed vector $\hat{\mathbf{k}}$ would not necessarily contain only kinematic information, and it would not necessarily be physically meaningful (i.e.: each regressed value corresponding to a different specific joint). For this reason, the regressor ψ is followed by a physical model of robotic instruments and surgical camera. This module ρ , called *physical module*, maps the estimated kinematics $\hat{\mathbf{k}}$ to the instruments 3-D shape, through a forward kinematic model of the instruments, and reprojects it to the camera plane as a binary mask \hat{m} , representing the instruments in the image




Fig. 3. *Backgroundizer module* β : the imprecise kinematics \mathbf{k}_s is converted to the binary projection representation m_s , through a robot-renderer model equivalent to ρ . m_s is then expanded to account for uncertainties and used to mask the image x associated with \mathbf{k}_s , which is then fed to an inpainting network which produces x_b , the *backgroundized* version of x .

space, according to the estimated kinematics $\hat{\mathbf{k}}$. The *physical module* ρ has the pivotal role of a *kinematic bottleneck*: because of the way it processes the low dimensional vector $\hat{\mathbf{k}}$, it gives it the explicit meaning of “kinematics”, forcing, in turn, ψ to learn the expected regression transformation, while keeping the self-supervised formulation of the problem. The outputs of the two branches, x_b and \hat{m} , are finally fed to the decoder ϕ , which tries to reconstruct x from the combination of the two. The *backgroundizer* and the *physical* modules are discussed in sections II-A and II-B.

A. Backgroundizer module: an inpainting problem

In order to make the reconstruction task feasible, the decoder must be provided with information about the appearance of the image. In [27], in human pose-estimation context, Jakob et al. provide this information by feeding a second image x' to the decoder, shifted in time with respect to x (with x, x' belonging to the same video). This approach is effective as long as x' contains the same exact background as x and the subject in it has a different pose than the one the model is trying to regress. Unfortunately, none of the hypotheses can be verified in a complex surgical scenario: the robotic instruments can remain in the same configuration for relatively long periods during a surgery, and the background appearance is constantly modified by the direct and indirect action of the instruments (pulling tissue, bleeding, smoke etc.) and by the movement of the camera. To address this issue, we introduce a novel *backgroundizer module* β that informs the decoder about the aspect of the image to reconstruct by producing the *backgroundized* image x_b as:

$$x_b = \beta(x, \mathbf{k}_s). \quad (2)$$

The image x_b is obtained by exploiting the rough instruments localization provided by \mathbf{k}_s , accounting for its uncertainty. In order to leverage the imprecise information provided by the input kinematics \mathbf{k}_s , we first map it to the corresponding binary mask m_s through a robot-renderer model, equivalent to ρ . We take into account the possible error in \mathbf{k}_s by expanding the binary mask m_s , in order to cover a greater space in the image. To do that, we simply consider the tool presence in quadrants obtained by overlapping a grid with $a \times b$ resolution on the mask m_s , and assigning to the whole quadrants a

value of 1 if at least one pixel of m_s inside is 1. The input image x and the expanded mask are then processed to obtain the *backgroundized* image x_b : the problem is formulated as image inpainting from block occlusion [28]. To solve it we implemented and trained a neural network model based on partial convolutions, following the implementation proposed in [29]. The model is trained using a perceptual loss function, as defined in [30]. The full pipeline for the *backgroundizer module* can be observed in fig. 3.

B. Physical module

The *physical module* ρ , following the regressor ψ , has the crucial role of *kinematic bottleneck*, giving $\hat{\mathbf{k}}$ the explicit meaning of “kinematics”, and forcing ψ to properly learn the kinematic regression task, while keeping the self-supervised formulation of the problem. This module consists of a geometrical, forward-kinematic model of the instruments, followed by a renderer (pinhole model of the surgical camera and rasterizer). The geometric forward kinematic model of the instruments is defined as the mapping $r : \mathbf{k} \rightarrow g$, with g being the 3-D shape of the instruments, referred to the instruments reference frame $\{r_{rf}\}$. Knowing the parameters of the rigid transformation T between $\{r_{rf}\}$ and the robot-mounted camera reference frame $\{c_{rf}\}$, and its calibration matrix, the instruments 3-D shape g can be projected on the camera plane and rasterized in pixel coordinates. The full rendering operation is defined as $c : g \rightarrow \hat{m}$, with \hat{m} being the binary instruments segmentation mask, obtained binarizing the projection of the 3-D shape of the instruments g on the camera plane.

The full problem can be so formulated as:

$$\hat{\theta}_\psi, \hat{\theta}_\phi = \arg \min_{\theta_\psi, \theta_\phi} \phi(\theta_\phi, \rho(\psi(\theta_\psi, x)), \beta(x, \mathbf{k}_s)) - x. \quad (3)$$

Given their embedding in an end-to-end trainable neural network architecture, both the instruments model r and the renderer c are required to be differentiable, in order to allow the error signal to be backpropagated through them. Recently, following the introduction by Kato et al. [31] of a differentiable 3-D mesh renderer, specifically designed for neural networks, some works have successfully introduced renderers in their pipelines, as in [32]. In this work, we address these challenges by directly learning the mapping between the kinematic vector $\hat{\mathbf{k}}$ and the output binary mask \hat{m} . Given a geometric forward kinematic model of the instruments and a renderer (both non-necessarily differentiable), one can generate a virtually infinite set of coupled samples $\{\mathbf{k}_i, m_i\}$ and use them to train a neural network to learn the direct mapping $\rho = r \circ c : \hat{\mathbf{k}} \rightarrow \hat{m}$. The network is implemented with a generator-like architecture, consisting of a first fully connected layer, whose output is reshaped and fed to a fully convolutional network, alternating normal convolution and deconvolution layers. The last layer of the network has a sigmoid activation function, and outputs an image \hat{m} , with same dimension as x and a single channel.

The training loss is formulated as a binary cross-entropy loss:

$$\mathcal{L}_j = -\frac{1}{N} \sum_{i=1}^{w \cdot h} \left[m_i^j \log(\hat{m}_i^j) + (1 - m_i^j) \log(1 - \hat{m}_i^j) \right], \quad (4)$$

with \hat{m}_i being the predicted value for pixel i , and w, h being the mask dimensions.

C. Regressor and Decoder

The regressor ψ processes the input image x to predict the corresponding instruments kinematic configuration $\hat{\mathbf{k}}$. Given the general independence between instruments (if more than one are simultaneously present), ψ is implemented using a shared backbone and m separate heads, one per robotic instrument. The backbone is based on ResNet50 [33], using only the first and second convolutional layers (21 residual blocks). Each head consists of four convolutional layers, followed by global average pooling and a three-layers fully connected network, having n/m units in output. In order to take advantage of the rough information provided by the measured kinematics k_s , without using it as a strong supervision signal, we introduce a “soft mean squared error” auxiliary loss function \mathcal{L}_a , defined as:

$$\mathcal{L}_a = \frac{1}{n} \sum_{i=1}^n \max((k_i - \hat{k}_i)^2 - t^2, 0), \quad (5)$$

with t being an hyperparameter, which we refer to as “tolerance”, and that can be interpreted as a maximum accepted offset of $\hat{\mathbf{k}}$ with respect to \mathbf{k}_s . The idea behind this loss is to ease the optimization process by providing the network with a range in which the solution is likely to be found, avoiding any hard-coded constraint.

The decoder ϕ processes the predicted projection mask \hat{m} and the *backgroundized* image x_b to reconstruct the image \hat{x} . The network is implemented as a *UU-net*, an extension of the well-established *U-net* architecture [34], having two separate contracting paths for \hat{m} and x_b , and an expanding path where corresponding features from the two contracting branches are concatenated. In order to learn the auto-encoder formulation defined in eq. 3, enforcing the reconstructed image \hat{x} to match the input image x , we implemented a perceptual loss \mathcal{L}_r , defined as:

$$\mathcal{L}_r = \|\Gamma(x) - \Gamma(\hat{x})\|^2, \quad (6)$$

where Γ is a feature extractor implemented as a VGG-16 network [35] pre-trained on the publicly available Imagenet dataset [36]. Compared to pixel-wise losses, the perceptual loss has shown more robust results in many image reconstruction tasks [37][38].

III. EXPERIMENTAL VALIDATION

The experimental validation was performed using the STRAS robot [10], a teleoperated prototype for flexible robotic endoscopic surgery [39]. The robot is built as a standard endoscope having two operating instrument channels, through which robotic arms can be positioned. The robot arms used in STRAS are flexible cable-actuated instruments. Each instrument has 3 joint angles/positions, resulting in a total dimension of $\hat{\mathbf{k}}$ equal to 6: rotation around the instrument main axis, translation along the same axis, and a cable-actuated bending, defined by the delta between the lengths of the two cables used for actuation. The instruments main axes are parallel to

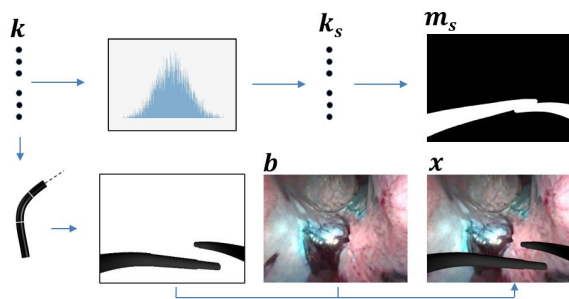


Fig. 4. Semi-synthetic dataset building: a kinematic configuration k and a background b are randomly sampled. k is then fed to a VTK model of the robotic instruments and rendered on the image space. The projection is then blended with b to obtain the image x . Parallely, a random noise is added to k to simulate the measured kinematics k_s (visualized as the corresponding binary projection m_s on the image plane, in order to qualitatively show the difference with the GT kinematic configuration k).

the exit parts of the instrument channels, which deviate from the endoscope main axis, forming a small angle ($\sim 10^\circ$), resulting almost perpendicular to the camera plane. During working configuration, instruments are bent, in order to be inside the field-of-view of the camera; therefore, when the rotational joint is activated, the visible part of the instruments moves on a plane almost parallel to the camera plane, avoiding ambiguities between pose and projected shape [11].

Using the constant curvature assumption, the joint angles can be used to compute the forward kinematic model of the robotic instruments, following equations detailed in [40]. The robot instruments are teleoperated by the user sending commands through a master console. The positions of the motors are recorded into the array of joint values k_s . Camera images are acquired through an acquisition board in a synchronized fashion, resulting in RGB images with 570×760 resolution. The geometric forward kinematic model of the instruments and the renderer were implemented using the VTK library. The two VTK models were used exclusively in the datasets generation process, as detailed later in this paragraph, and not included in the training architecture, given their non-differentiability.

Three different datasets were used for validating the proposed approach, two from real acquisitions and a semi-synthetic one. The first real dataset (*phantom*) was recorded on the benchtop, using a plastic phantom model of the human digestive system. During the acquisition both the endoscope and the plastic model were moved, in order to avoid having a static background appearance. The second real dataset (*in-vivo*) was recorded during a 4-day set of *in-vivo* experiments on porcine models¹. This dataset presents several challenges compared to standard retinal, laparoscopic and benchtop datasets, including low foreground/background contrast, highly cluttered and changing background, frequent occlusions and bleedings, and strong tool-tissue interactions. The *semi-synthetic* dataset was created by blending background-only images, automatically

¹The study protocol for this experiment was approved by the Institutional Ethical Committee on Animal Experimentation (ICOMETH No.38.2011.01.018). Animals were managed in accordance with French laws for animal use and care as well as with the European Community Council directive no. 2010/63/EU

TABLE I
NUMBER OF IMAGES IN EACH TRAINING AND TESTING DATASET (* GT 2-D SEGMENTATION MASK ONLY, OBTAINED VIA MANUAL SEGMENTATION)

Dataset	# Images Training	# GT testing images
<i>physical module</i>	100k	10k
<i>real backgrounds</i>	6000	/
<i>phantom backgrounds</i>	1 (+augmentation)	/
<i>semi-synthetic</i>	20400	2400
<i>phantom</i>	6800	800*
<i>in-vivo</i>	28800 (4 days)	400*/day

extracted from the *in-vivo* dataset by parsing the associated kinematic information (although imprecise), with rendered robot instruments, obtained using the VTK model of tools and renderer, and random kinematic configurations k . In order to simulate a realistic imprecise kinematic information k_s for this dataset, we added to the nominal value k a normally distributed noise, whose range was empirically chosen to match the real-dataset noise. The full process followed to build the semi-synthetic dataset is shown in fig. 4.

For the *semi-synthetic* dataset, the ground truth (GT) joint configuration is known a-priori, due to the way the dataset was built. For the *phantom* and *in-vivo* datasets, the GT joint configuration is unknown, since every measurement attempt is affected by uncertainties, resulting on the imprecise estimation k_s . Instead, the evaluation on these two datasets was performed indirectly: the 3-D shape of the instruments is reconstructed according to the predicted kinematics \hat{k} , and projected on the camera plane: the projected mask \hat{m} can be then compared to the instruments ground truth location on the images, obtained via manual segmentation, providing an indirect evaluation of the estimated \hat{k} .

In order to train and evaluate the *physical module*, a dataset was built generating random couples $\{k_i, m_i\}$, using the VTK model of instruments and renderer. For the training of the inpainting network, two separate datasets were built. The first one, *real backgrounds*, was built automatically extracting background-only images from the *in-vivo* dataset, according to the associated kinematic information (although imprecise), and used to train the *backgroundizer* for the *semi-synthetic* and *in-vivo* experiments; the second one, *phantom backgrounds* was built by extracting a single background-only image from the *phantom dataset*, strongly augmented through operations like rotation, cropping, lighting etc., and used for the *phantom* experiment. Table I summarizes the number of images and GT images for the datasets.

A. Training Details

As a preliminary step, the inpainting network, belonging to the *backgroundizer module* β , and the neural network implementation of the *physical module* ρ were trained. The *physical module* was trained on the randomly generated couples $\{k_i, m_i\}$, for 100 epochs, using a batch size of 64 and a learning rate of 0.0005. For the inpainting network the grid resolution was set to 6×8 , as experimentally found providing the best trade-off between area covered and inpainting quality. Two inpainting networks were trained (on *real backgrounds*

TABLE II

SEMI-SYNTHETIC DATASET RESULTS. COMPARISON WITH RAW KINEMATICS k_s AND FULLY SUPERVISED METHODS $BSupK_s$, $BSupSoftK_s$. FOR THE JOINT MEAN ABSOLUTE ERRORS (TRANSLATION: t ., ROTATION: r ., BENDING: b .), LOWER IS BETTER. FOR THE REPROJECTION METRICS (INTERSECTION OVER UNION: IoU), HIGHER IS BETTER.

	<i>left</i>				<i>right</i>			
	t .[mm]	r .[deg]	b .[mm]	IoU	t .[mm]	r .[deg]	b .[mm]	IoU
k_s	6.40	16.40	0.87	0.244	3.21	13.07	0.77	0.422
$BSupK_s$	4.00	10.30	0.73	0.646	1.85	7.05	0.49	0.646
$BSupSoftK_s$	4.27	11.00	0.62	0.355	2.13	6.97	0.50	0.620
<i>ours</i>	1.75	6.02	0.47	0.738	1.17	3.61	0.30	0.854

TABLE III

EVALUATION OF THE IOU ON THE REAL DATASETS (PHANTOM & IN-VIVO). COMPARISON WITH RAW KINEMATICS k_s AND FULLY SUPERVISED METHODS $BSupK_s$, $BSupSoftK_s$. FOR THE IN-VIVO THE AVERAGE OF THE RESULTS OBTAINED FOR EACH DAY IS REPORTED.

	<i>left</i>		<i>right</i>	
	<i>phant.</i>	<i>in-vivo</i>	<i>phant.</i>	<i>in-vivo</i>
k_s	0.283	0.421	0.327	0.447
$BSupK_s$	0.280	0.482	0.327	0.461
$BSupSoftK_s$	0.313	0.458	0.330	0.436
<i>ours</i>	0.640	0.554	0.725	0.554

and *phantom backgrounds*), with a batch size of 12 and a learning rate of 0.001, until visually satisfying results were reached.

The two modules were then frozen during the end-to-end training of the regressor ψ and the decoder ϕ . The regressor and the decoder were trained alternatively (1 iteration each) using the Adam optimizer with learning rates equal to 0.0005 and 0.00001, respectively, a batch size of 36, for 100 epochs: the decoder is trained on loss \mathcal{L}_r ; the regressor is trained on a weighted sum of losses \mathcal{L}_r and \mathcal{L}_a , defined as:

$$\mathcal{L} = \alpha_a \mathcal{L}_a + \alpha_r \mathcal{L}_r, \quad (7)$$

with α_a , α_r being weight parameters set to 10 and 0.001, respectively, in order to balance the magnitude of the two losses. The tolerance parameter t was tuned through an hyperparameter search, adjusting a nominal value empirically determined in the context of other works using the STRAS robotic system [41][10].

IV. RESULTS

In this section we present the validation results obtained by the proposed approach on the three datasets. In order to assess the actual contribution of the method, we compared its performances to two baseline models. They both consist of solely the regressor model ψ trained by minimizing only the loss function \mathcal{L}_a defined in eq. 5. The first baseline model $BSupK_s$ was trained by setting the tolerance parameter t to 0, thus resulting in a strongly-supervised training, having k_s , which is unreliable/imperfect, as target; the second baseline model $BSupSoftK_s$ was trained by setting the parameter t to the same value as the one used for our main model, thus resulting in a softer supervision. Together with the baselines, also the

raw measured kinematics k_s provided by the robotic system was evaluated, according to the same modalities.

For the evaluation on the *in-vivo* dataset, leave-one-out cross-validation was performed, by training the models on 3-days and testing on the remaining. The average of the results across the 4 days was then computed and reported.

A. Backgroundizer and physical modules preliminary evaluation

The *physical module*, preliminary trained, was evaluated resulting in a mean Intersection over Union (IoU) of 94.3% on the testing dataset. The inpainting networks (*backgroundizer* module), preliminary trained, were qualitatively evaluated. Qualitative results for the three datasets are shown in fig. 6.

B. Regressor evaluation

For the *semi-synthetic* dataset, the GT kinematic configuration k of the instruments is known, and the mean absolute error for each joint could be computed. Reprojection of the estimated instruments shape on the image plane was also evaluated through the standard segmentation metric Intersection over Union (IoU). Results are reported in table II. For the real *phantom* and *in-vivo* datasets the reprojection error with respect to the manually annotated GT was evaluated, via IoU. Results are reported in table III. Qualitative results for the three datasets can be observed in fig. 5.

Results on the three datasets, despite not reaching a state-of-the-art accuracy, show a consistent improvement provided by the proposed approach. This is shown directly by the improved kinematics estimation in the *semi-synthetic* dataset, and indirectly by the higher reprojection accuracies on the *semi-synthetic* and real datasets, with respect to both the baselines and the raw kinematics. In the *semi-synthetic* and *in-vivo* datasets the baselines perform better than the raw kinematics k_s , while in the *phantom* dataset the performances are comparable. This may be due to the smaller size of the *phantom* dataset, which prevents the networks from learning to average out the random noise of k_s with respect to the GT value k . Our model, on the contrary, is not affected, proving its capability to learn also from reduced amount of noisy data. The evaluation of $BSupK_s$ provides an ablation study of our method, showing that the improvements reached are not solely caused by the soft loss. In the *in-vivo* dataset, a lower improvement is reached. This can be explained by two main reasons: 1) the dataset is more challenging due




Fig. 5. Examples of qualitative results on the three testing datasets. Top row: *semi-synthetic*; middle row: *phantom*; bottom row: *in-vivo*. The last two columns show corresponding reconstructed image and *backgroundized* image. Note that at inference time none of the two is needed/obtained, since the image-based regressor ψ is a completely independent module.




Fig. 6. Examples of *backgroundizer* module results. Top row shows one image for each dataset (*semi-synthetic*, *phantom*, *in-vivo*), with the corresponding measured kinematics k_s projection (imprecise). Bottom row shows the corresponding *backgroundized* versions, obtained using the rough localization provided by k_s .

to the factors mentioned in section III (e.g. low instruments-background contrast), which can affect our model more than the baselines, since the training is almost completely image-based; 2) phenomena as tool-channel interaction and slackening [42], more evident under strong tool-tissue interaction, can deform the flexible instruments in configurations that our model, based on constant curvature assumption, cannot reach. An additional interesting aspect of the proposed method is that the fully trained architecture can be divided in autonomous, fully-functional sub-modules:

- the regressor ψ , the core element of the method,
- a segmentation module, combination of ψ and ρ ,
- an image synthesizer ϕ , able to generate a real-looking sample from a background image and a GT segmentation mask of the tools (fig. 5, last column).

Finally, the average processing time for each image is approximately 30 ms (~ 33 fps) on a single Tesla V100 GPU.

V. CONCLUSION AND FUTURE WORK

In this work, we propose a method for 3-D pose estimation of flexible surgical instruments. The proposed approach is self-supervised and trains a model to regress the kinematic

configuration of the instruments directly from surgical camera images. In order to train this model with no manual annotations, the problem is formulated in an auto-encoder framework, by exploiting a tight *kinematic bottleneck* to force a regressor to learn the actual kinematics. Our approach also takes advantage, at training time only, of imprecise kinematic information provided by the robotic system along with the camera images. While the kinematic data is imprecise due to tool-tissue interactions and is usually difficult to access in real-time applications inside the operating room, it helps drive the learning process to a strong regressor that estimates even better kinematics. One strong advantage is that no manual annotation is required during training and no additional optimization is performed at inference time. Furthermore, in addition to the regressed kinematics, the model directly provides a segmentation mask of the tools.

Experimental evaluation was carried out using a particular robotic endoscope having two flexible robotic arms, each one characterized by three joint values, on three datasets: *semi-synthetic*, *phantom* and *in-vivo*. Results show a consistent improvement with respect to the baselines trained directly on the raw instruments kinematics.

While the results are promising, the proposed approach presents a few limitations that could be addressed. First, the kinematic model of the instruments should be improved, in order to take into account phenomena as tool-channel interactions and slackening, particularly evident in non-controlled environments, as the *in-vivo* dataset. Secondly, the uncertainty associated with the robot kinematics could be explicitly modelled, instead of relying on the naive grid-approach, with possible benefits on the *backgroundizer* and the whole training process. Finally, temporal dependencies between consecutive images could be modelled.

ACKNOWLEDGMENT

Thanks to Bernard Dallemagne, Florent Nageotte, and Philippe Zanne for the help in generating the datasets.

REFERENCES

- [1] J. H. Palep, "Robotic assisted minimally invasive surgery," *Journal of minimal access surgery*, vol. 5, no. 1, p. 1, 2009.
- [2] L. Maier-Hein, S. S. Vedula, S. Speidel, N. Navab, R. Kikinis, A. Park, M. Eisenmann, H. Feussner, G. Forestier, S. Giannarou, *et al.*, "Surgical data science for next-generation interventions," *Nature Biomedical Engineering*, vol. 1, no. 9, pp. 691–696, 2017.
- [3] B. Münzer, K. Schoeffmann, and L. Böszörményi, "Content-based processing and analysis of endoscopic images and videos: A survey," *Multimedia Tools and Applications*, vol. 77, no. 1, pp. 1323–1362, 2018.
- [4] H. Azimian, R. V. Patel, and M. D. Naish, "On constrained manipulation in robotics-assisted minimally invasive surgery," in *2010 3rd IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechanics*. IEEE, 2010, pp. 650–655.
- [5] O. M. Omisore, S. Han, Y. Al-Handarish, W. Du, W. Duan, T. O. Akinyemi, and L. Wang, "Motion and trajectory constraints control modeling for flexible surgical robotic systems," *Micromachines*, vol. 11, no. 4, p. 386, 2020.
- [6] A. M. Okamura, "Methods for haptic feedback in teleoperated robot-assisted surgery," *Industrial Robot: An International Journal*, 2004.
- [7] A. Talasaz and R. V. Patel, "Remote palpation to localize tumors in robot-assisted minimally invasive approach," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 3719–3724.
- [8] J. Burgner-Kahrs, D. C. Rucker, and H. Choset, "Continuum robots for medical applications: A survey," *IEEE Transactions on Robotics*, vol. 31, no. 6, pp. 1261–1280, 2015.
- [9] C. Shi, X. Luo, P. Qi, T. Li, S. Song, Z. Najdovski, T. Fukuda, and H. Ren, "Shape sensing techniques for continuum robots in minimally invasive surgery: A survey," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 8, pp. 1665–1678, 2016.
- [10] A. De Donno, L. Zorn, P. Zanne, F. Nageotte, and M. de Mathelin, "Introducing stras: A new flexible robotic system for minimally invasive surgery," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 1213–1220.
- [11] M. Allan, S. Ourselin, D. J. Hawkes, J. D. Kelly, and D. Stoyanov, "3-d pose estimation of articulated instruments in robotic minimally invasive surgery," *IEEE transactions on medical imaging*, vol. 37, no. 5, pp. 1204–1213, 2018.
- [12] M. K. Chmarra, C. Grimbergen, and J. Dankelman, "Systems for tracking minimally invasive surgical instruments," *Minimally Invasive Therapy & Allied Technologies*, vol. 16, no. 6, pp. 328–340, 2007.
- [13] A. Gadwe and H. Ren, "Real-time 6dof pose estimation of endoscopic instruments using printable markers," *IEEE Sensors Journal*, vol. 19, no. 6, pp. 2338–2346, 2018.
- [14] Y.-S. Wang and K.-T. Song, "Image-based pose estimation and tracking of surgical instruments in minimally invasive surgery," in *2020 International Automatic Control Conference (CACSC)*. IEEE, 2020, pp. 1–6.
- [15] D. Bouget, M. Allan, D. Stoyanov, and P. Jannin, "Vision-based and marker-less surgical tool detection and tracking: a review of the literature," *Medical image analysis*, vol. 35, pp. 633–654, 2017.
- [16] M. Allan, P.-L. Chang, S. Ourselin, D. J. Hawkes, A. Sridhar, J. Kelly, and D. Stoyanov, "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, 2015, pp. 331–338.
- [17] I. Laina, N. Rieke, C. Ruppert, J. P. Vizcaíno, A. Eslami, F. Tombari, and N. Navab, "Concurrent segmentation and localization for tracking of surgical instruments," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2017, pp. 664–672.
- [18] T. Kurmann, P. M. Neila, X. Du, P. Fua, D. Stoyanov, S. Wolf, and R. Sznitman, "Simultaneous recognition and pose estimation of instruments in minimally invasive surgery," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2017, pp. 505–513.
- [19] X. Du, T. Kurmann, P.-L. Chang, M. Allan, S. Ourselin, R. Sznitman, J. D. Kelly, and D. Stoyanov, "Articulated multi-instrument 2-d pose estimation using fully convolutional networks," *IEEE transactions on medical imaging*, vol. 37, no. 5, pp. 1276–1287, 2018.
- [20] J. Mu, W. Qiu, G. D. Hager, and A. L. Yuille, "Learning from synthetic animals," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 12 386–12 395.
- [21] T. Heimann, P. Mountney, M. John, and R. Ionasec, "Real-time ultrasound transducer localization in fluoroscopy images by transfer learning from synthetic training data," *Medical image analysis*, vol. 18, no. 8, pp. 1320–1328, 2014.
- [22] D. Pakhomov, W. Shen, and N. Navab, "Towards unsupervised learning for instrument segmentation in robotic surgery with cycle-consistent adversarial networks," *arXiv preprint arXiv:2007.04505*, 2020.
- [23] T. Yu, D. Mutter, J. Marescaux, and N. Padoy, "Learning from a tiny dataset of manual annotations: a teacher/student approach for surgical phase recognition," *arXiv preprint arXiv:1812.00033*, 2018.
- [24] T. Ross, D. Zimmerer, A. Vemuri, F. Isensee, M. Wiesenfarth, S. Bodenstedt, F. Both, P. Kessler, M. Wagner, B. Müller, *et al.*, "Exploiting the potential of unlabeled endoscopic video data with self-supervised learning," *International journal of computer assisted radiology and surgery*, vol. 13, no. 6, pp. 925–933, 2018.
- [25] C. I. Nwoye, D. Mutter, J. Marescaux, and N. Padoy, "Weakly supervised convolutional lstm approach for tool tracking in laparoscopic videos," *International journal of computer assisted radiology and surgery*, vol. 14, no. 6, pp. 1059–1067, 2019.
- [26] A. K. Tanwani, P. Sermanet, A. Yan, R. Anand, M. Phielipp, and K. Goldberg, "Motion2vec: Semi-supervised representation learning from surgical videos," *arXiv preprint arXiv:2006.00545*, 2020.
- [27] T. Jakob, A. Gupta, H. Bilen, and A. Vedaldi, "Self-supervised learning of interpretable keypoints from unlabelled videos," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 8787–8797.
- [28] O. Elharrouss, N. Almaadeed, S. Al-Maadeed, and Y. Akbari, "Image inpainting: A review," *Neural Processing Letters*, pp. 1–22, 2019.
- [29] G. Liu, F. A. Reda, K. J. Shih, T.-C. Wang, A. Tao, and B. Catanzaro, "Image inpainting for irregular holes using partial convolutions," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 85–100.
- [30] L. A. Gatys, A. S. Ecker, and M. Bethge, "A neural algorithm of artistic style," *arXiv preprint arXiv:1508.06576*, 2015.
- [31] H. Kato, Y. Ushiku, and T. Harada, "Neural 3d mesh renderer," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3907–3916.
- [32] S. Wu, C. Ruppert, and A. Vedaldi, "Unsupervised learning of probably symmetric deformable 3d objects from images in the wild," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1–10.
- [33] M. Boroumand, M. Chen, and J. Fridrich, "Deep residual network for steganalysis of digital images," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 5, pp. 1181–1193, 2018.
- [34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [35] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [36] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [37] J. Bruna, P. Sprechmann, and Y. LeCun, "Super-resolution with deep convolutional sufficient statistics," *arXiv preprint arXiv:1511.05666*, 2015.
- [38] A. Dosovitskiy and T. Brox, "Generating images with perceptual similarity metrics based on deep networks," in *Advances in neural information processing systems*, 2016, pp. 658–666.
- [39] A. Légner, M. Diana, P. Halvax, Y.-Y. Liu, L. Zorn, P. Zanne, F. Nageotte, M. De Mathelin, B. Dallemagne, and J. Marescaux, "Endoluminal surgical triangulation 2.0: A new flexible surgical robot. preliminary pre-clinical results with colonic submucosal dissection," *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 13, no. 3, p. e1819, 2017.
- [40] R. J. Webster III and B. A. Jones, "Design and kinematic modeling of constant curvature continuum robots: A review," *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1661–1683, 2010.
- [41] C. da Costa Rocha, N. Padoy, and B. Rosa, "Self-supervised surgical tool segmentation using kinematic information," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8720–8726.
- [42] D. B. Camarillo, C. F. Milne, C. R. Carlson, M. R. Zinn, and J. K. Salisbury, "Mechanics modeling of tendon-driven continuum manipulators," *IEEE transactions on robotics*, vol. 24, no. 6, pp. 1262–1273, 2008.