

© 20XX IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Deliberation in autonomous robotic surgery: a framework for handling anatomical uncertainty

Eleonora Tagliabue*, Daniele Meli*, Diego Dall’Alba, Paolo Fiorini

Abstract—Autonomous robotic surgery requires deliberation, i.e. the ability to plan and execute a task adapting to uncertain and dynamic environments. Uncertainty in the surgical domain is mainly related to the partial pre-operative knowledge about patient-specific anatomical properties. In this paper, we introduce a logic-based framework for surgical tasks with deliberative functions of monitoring and learning. The Deliberative Framework for Robot-Assisted Surgery (DEFRAS) estimates a pre-operative patient-specific plan, and executes it while continuously measuring the applied force obtained from a biomechanical pre-operative model. Monitoring module compares this model with the actual situation reconstructed from sensors. In case of significant mismatch, the learning module is invoked to update the model, thus improving the estimate of the exerted force. DEFRAS is validated both in simulated and real environment with da Vinci Research Kit executing soft tissue retraction. Compared with state-of-the-art related works, the success rate of the task is improved while minimizing the interaction with the tissue to prevent unintentional damage.

Index Terms—Planning under uncertainty; Surgical robotics: planning; Software architectures for robotics

I. INTRODUCTION

In recent years, research in surgical robotics has focused on the automation of surgical tasks, for its potential to provide physical and cognitive support to the surgeon [1]. Existing works in this field have mainly considered structured and controlled environments [2]. This assumption makes current approaches still far from the applicability to real clinical scenarios, which consist of deformable tissues whose mechanical properties are unknown pre-operatively, leading to intra-operative uncertainty [3].

To cope with real deformable anatomical environments, an Autonomous Robotic Surgery (ARS) system shall integrate the following skills:

- it should have *deliberative* capabilities [4], i.e. *monitoring* the environment through *sensors* while *acting*, in order to *reason* and *plan* a new strategy when the expected behavior does not match reality. Moreover, it must adapt its prior model by *learning* from observations, to minimize uncertainty and critical events;
- it should be *informed* from pre-operative anatomical data and task-level expert knowledge, in order to devise a patient-specific strategy for the intervention;

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 742671 (ARS).

*Authors contributed equally to the work.

Authors are with Dept. of Computer Science, University of Verona, Italy. Contact author: Daniele Meli, daniele.meli@univr.it

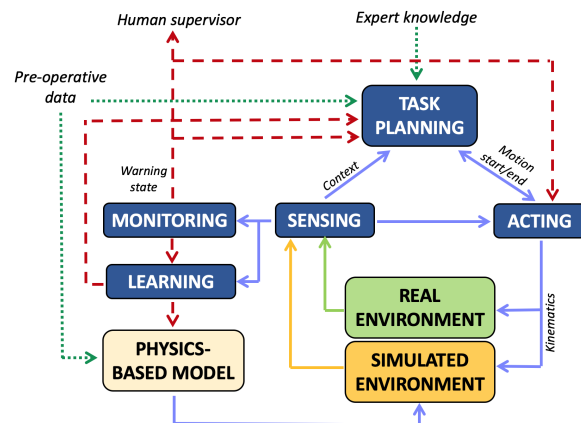


Fig. 1. DEFRAS architecture (deliberative functions in blue boxes). Green lines represent pre-operative information. Red lines trigger actions only when monitoring threshold is exceeded.

- it should provide *interpretable* plans, which can be easily understood by a supervising surgeon, to guarantee reliability of the ARS system [3], [5], as required from level 2 of autonomy [1].

To address autonomy in real deformable surgical environments, we rely on a biomechanical simulation initialized with pre-operative information both to generate a patient-specific surgical plan *pre-operatively* and to detect possible unexpected behavior *intra-operatively*, via comparison with real sensor data. The specific contributions of this work are:

- 1) we present DEFRAS, a DELiberative Framework for Robot-Assisted Surgery that integrates skills necessary to cope with the uncertainties and the dynamic behavior of anatomical environments, while guaranteeing task completion and respect of safety constraints;
- 2) we demonstrate it for a representative surgical task with uncertainty, i.e. soft tissue retraction, both in simulated and real experiments with the da Vinci Research Kit (dVRK) [6];
- 3) we publicly share the framework at https://gitlab.com/altairLab/tissue_retraction.git.

The paper is organized as follows. In Section III we review recent frameworks for surgical robotic autonomy under uncertainty and generic frameworks for deliberation; we then describe DEFRAS and our implementation of the tissue retraction task in Sections III and IV respectively; Section V describes the conducted validation experiments and Section VI reports the obtained results. Finally, we draw our conclusions in Section VII.

II. RELATED WORKS

Uncertainty poses challenges to ARS systems relying only on prior knowledge about anatomical models and/or task description [7]. Previous works have tried to deal with the different sources of uncertainty in the surgical domain, e.g. implementing strategies for accurate instrument and anatomy localization [8], [9], online update of soft tissues models [10], [11], workflow recognition [12] and task model refinement [13], [14]. Model predictive control has also been exploited to deal with uncertainty at motion planning level [12], [15].

While existing works have attempted either to reduce uncertainty at sensing level or to deal with it at motion/task level, in this paper we address both aspects by introducing deliberative functions of autonomous agents [4]. In particular, we integrate specific modules for uncertainty *monitoring* and pre-operative model refinement through online *learning* within a framework for robot-assisted surgery [16] that guarantees task-level interpretability with logic programming and supports online re-planning. Several paradigms for deliberation have been proposed in the field of artificial intelligence, focusing on different aspects of autonomy, e.g. planning [17]–[19] and learning [20]. However, in order to guarantee adaptability to diverse complex tasks, here we focus on the integration and parallel development of different skills for autonomy, which can be improved individually depending on the task of interest. For this reason, our framework builds on the modular architecture and definitions of deliberative functions described in [4], which reports a large number of robotic applications.

III. DEFRAS: DELIBERATIVE FRAMEWORK FOR ROBOT-ASSISTED SURGERY

Our deliberative architecture is shown in Fig. 1. A *task planning* module computes the task plan, i.e. the sequence of actions to solve the task. Actions are elementary operations which correspond to a single motion trajectory or joint-level command (e.g., reaching a target or grasping), according to the granularity definitions described in [21]. The plan is computed based on available expert procedural knowledge, conveniently expressed in the declarative logic programming formalism of Answer Set Programming (ASP) [22], [23]. The ASP task description encodes entities, i.e. definitions of actions, tools and relevant anatomical features as logic variables and predicates (*atoms*). Logical implications between atoms are implemented to define pre-conditions and effects of actions, and constraints (e.g., infeasibility of some actions). *Answer sets* are then computed by ASP solver Clingo [24] as the minimal sets of grounded atoms (i.e., with assignments of variables), hence including the shortest plans along with the corresponding environmental features. This enhances task-level interpretability. Throughout task execution, the *sensing* module analyses the state of the robot and the environment, to provide high-level environmental features for task planning and targets for motion planning (*acting* module). Moreover, it recognizes whether critical conditions occur, in which cases it triggers task re-planning or motion interruption to prevent risks. The *monitoring*

module continuously estimates if the internal model of the environment, given by a patient-specific simulation initialized with pre-operative information and running in parallel, is coherently representing the actual scenario. To this aim, a task-specific metric assessing the level of discrepancy between the real and the simulated environments is computed starting from sensing information. Depending on discrepancy, different *warning states* are raised, associated with specific strategies at task-motion level to increase reliability and accommodate safety. Finally, the *learning* module is in charge of updating the pre-operative model from real data in warning states, to compensate for environmental uncertainty caused by inaccurate model parametrization.

IV. TISSUE RETRACTION WITH DEFRAS

DEFRAS functionalities are tested on autonomous execution of Tissue Retraction (TR), a representative surgical task which involves the interaction with soft anatomical tissues, whose properties might be not fully known before the intervention [25], [26]. It consists in grasping and retracting soft tissues to expose a hidden Region Of Interest (ROI) to an intra-operative camera (Fig. 3). The tissue is attached to surrounding anatomies in correspondence of Attachment Points (APs). Previous works for autonomous TR have proposed either to generate the plan by optimizing some metrics [26]–[30] or to learn the task from human demonstrations [25], [31], [32]. However, none of them integrates methods to reason on the current environment and deal with unexpected situations which might occur in real surgery.

A. TR implementation with DEFRAS

The task is executed on the dVRK, whose Patient-Side Manipulators (PSMs) are equipped with ProGrasp® surgical grasping tools. In general, ROI exposure is obtained primarily by pulling up the tissue after grasping. The grasping location on the tissue can be reached by either one of the PSMs, depending on *reachability* conditions defined with respect of the tissue center¹ to avoid instrument collision during operation. However, pulling is not always successful due to kinematic limitations imposed by the anatomical environment and the insertion points of the instruments in the body, which define fixed pivots for the robotic arms. Furthermore, pulling may induce high forces which can damage the tissue. In these cases, alternative actions can be performed, such as moving the tissue in a parallel plane to the tissue surface, either folding it away from the intra-operative camera or possibly towards the ROI.

In this work, we create a physics-based model of the anatomy leveraging on tissue geometric and mechanical properties as well as locations of the ROI and Attachment Points (APs) obtained from pre-operative data. Such model is used to initialize a simulation which is exploited to find a patient-specific task plan in the pre-operative phase. We select the plan that maximizes the economy of motion and reduces the risk of critical events (see Section V for

¹With reference to Fig. 3 a point on the tissue is reachable by a PSM if it lies on the same side as the PSM with respect to the $x - z$ plane.

details). The pre-operative plan is then executed in the real environment.

During execution, the physics-based simulation is run in parallel, and the state of the real environment is continuously monitored with intra-operative sensors (Fig. 1). When mismatch between simulation and reality is identified, the pre-operative plan may become sub-optimal or even infeasible, hence simulation parameters may be updated and specific strategies to improve the plan are implemented. In this work, we focus on the update of APs. Although APs can be initialized pre-operatively based on anatomical knowledge, previous works have shown that their proper definition is possible only during interaction with the tissue [33], [34]. The reason why we focus on APs is twofold. Firstly, they constitute model boundary conditions, thus represent an important parameter influencing simulation accuracy [33], [34]. Secondly, knowledge of APs location is very important in TR to prevent the pulling of highly constrained areas, minimizing the risk of tissue damage. The online simulation also continuously provides an estimate of the applied force on the tissue, compensating for the lack of force feedback in surgical robotic systems.

B. Task planning

Possible actions implemented in ASP task description are *reaching a grasping point*, *opening gripper*, *grasping*, *pulling* the tissue and *folding* it, either *moving planarly towards the ROI* or *moving away from the intra-operative camera*. They are expressed as predicates in plain English *action(tool, object, property, time)*, e.g., `reach(A, point, (X, Y), t)`, being A a generic tool, (X, Y) the coordinates of a grasping point and t a discrete time step. Possible grasping points are defined as ASP entities to be evenly spaced on the tissue surface. This allows to exploit built-in ASP optimization to select the most reasonable grasping point in case of re-planning, as the one farthest from APs and closest to ROI [16]. This is a commonsense approach to minimize tissue damage and expose ROI fast, but different strategies are also possible, e.g., [28]. Relevant environmental features include *reachability* conditions, *mutual locations of APs, ROI and tools* and conditions on *force* and *height* of PSMs, as in [16]. They are expressed as predicates similarly to actions (e.g., `max_force(A, t)`). Logical implications encode action pre-conditions, effects, constraints and the goal (i.e., the ROI must be eventually exposed). For instance, *grasp only reachable points* becomes `grasp(A, point, (X, Y), t) :- reachable(A, point, (X, Y), t)`, where `:-` is ASP notation for \leftarrow .

C. Sensing and acting

The sensing module continuously receives the state of the robot and the environment. The state of the robot consists of the position of tools and opening angles of grippers. The state of the environment includes the point cloud of tissue surface PC_t and the ROI from the intra-operative camera; the position of the ROI, known APs and estimated

Algorithm 1 Monitoring module

```

1: Input:  $PC_t, PC_s$ , motion rate  $r$ , grasped.tissue
2: Output: Warning state  $s$ , motion rate  $r$ , trigger learning
   learn
3: Init:  $d_{rest} = []$ 
4: while not task_ended do
5:   while not grasped.tissue do
6:      $d_{rest}.append(\|PC_t - PC_s\|_2)$ 
7:    $d_{rest} = \text{mean}(d_{rest})$ 
8:    $d = \|PC_t - PC_s\|_2$ 
9:    $\mu = d - d_{rest}$ 
10:  if  $\delta_1 < \mu \leq \delta_2$  then
11:    return  $s = 1; r = \frac{r}{2}; \text{learn} = \text{True}$ 
12:  else if  $\delta_2 < \mu \leq \delta_3$  then
13:    return  $s = 2; r = 0; \text{learn} = \text{True}$ 
14:  else if  $\mu > \delta_3$  then
15:    return  $s = 3; r = 0; \text{learn} = \text{False}$ 
16:  else
17:    return  $s = 0; r = r; \text{learn} = \text{False}$ 

```

forces at each tissue point from the simulation. Geometric reasoning is performed in real time on input data to ground anatomical features for task planning and anomaly detection. In particular, the sensing module continuously checks whether excessive force is applied to the tissue. When the maximum tissue force F exceeds a pre-defined threshold ϵ , DEFRAS first tries to reduce tissue solicitation by decreasing motion velocity [35], otherwise motion is interrupted and re-planning is triggered. Input data is also used to compute target positions for actions (e.g., locations of grasping points for *reach*). Orientations of PSMs are kept constant for simplicity. The acting module then translates ASP actions to joint (*grasp, release*) or Cartesian trajectories for the dVRK.

D. Monitoring

The monitoring module continuously receives PC_t and the point cloud of the simulated tissue PC_s from the sensing module, and computes a discrepancy metric μ , chosen as the median Euclidean distance of closest points between PC_t and PC_s (Alg. 1). The distance is normalized with respect to its rest value (i.e., before grasping occurs) to compensate for offsets induced by registration error and sensor noise. Inspired from recent works as [36], the monitoring module considers 3 thresholds $\delta_1, \delta_2, \delta_3$ (empirically set to 2, 4 and 6 mm) to classify the system state. In particular, when $\mu > \delta_1$ the motion velocity is reduced for caution, to minimize the risk of applying excessive forces [35] and allow prompt response to anomalies. If δ_2 is exceeded, motion is interrupted and re-planning is asked at task level. In both states, the learning module is triggered to update environmental knowledge exploiting data from sensors. Overcoming δ_3 results in interruption of DEFRAS and human intervention is required.

E. Learning

The learning module focuses on the intra-operative update of APs whenever required by the monitoring module. Update relies on BANet, a deep neural network which predicts APs starting from the tissue pre-operative geometry and the current observed PC_t [37]. This method has proved suitable for intra-operative model refinement, requiring a very low inference time and even improving biomechanical model accuracy [34]. The new estimate of APs is used both to refine the simulated physics-based model, thus improving its reliability (especially of the force estimate), and to update the knowledge of the regions to avoid in case of re-planning.

V. EXPERIMENTS

We consider the problem of retracting a thin layer of soft tissue (120x120x5 mm) to expose at least 80% of ROI surface to the endoscopic camera, which is placed at the main surgical viewpoint (Fig. 3). In particular, we aim at assessing the capability of DEFRAS to accomplish TR task, coping with inaccurate parametrization of the available model of the environment from pre-operative information.

For monitoring purposes, we continuously capture the tissue point cloud PC_t with an additional Intel® Realsense D435 RGBD sensor, placed opposite to the endoscope to maximize tissue visibility during manipulation, as done in [34]. The RGBD sensor, endoscopic camera and the PSMs are calibrated with respect to a common reference frame placed at the center of the tissue, using the methodology described in [38]. The simulated environment is aligned to the same reference frame and a simulated RGBD sensor placed in the same position as the real one provides the point cloud of the simulated tissue PC_s . Tissue deformable behavior is modelled relying on the laws of continuum mechanics, solved with the finite element method using the SOFA framework [39]. The tissue is described as a linear elastic material with Young’s modulus 3kPa and Poisson ratio 0.45 [34], relying on the corotational formulation to handle large deformations while guaranteeing computational efficiency. For all experiments in this paper, we consider a set of 25 possible grasping points, regularly spaced over a 5x5 grid to maximize computational efficiency [16]. We design 17 pre-operative configurations of tissue with different locations for APs and ROI (Fig. 2a-3b).

As an initial step, we perform an exploratory search to identify a reasonable value for the force threshold ϵ . For each pre-operative configuration we launch DEFRAS only in simulation, deactivating learning and monitoring and ignoring force measurements. We compute plans for all possible grasping points under these conditions, and we record the estimated forces from the simulation. We obtain a median force $F_m = 0.15$ N among all 425 executions. F_m is chosen to represent the typical force on the tissue in the considered setup, hence we set $\epsilon = 2F_m = 0.3$ N.

Afterwards, for each pre-operative configuration we generate an optimal patient-specific TR plan as follows. We launch DEFRAS in simulation without learning and monitoring modules, but now including force measurements. We record

plans originated picking all possible grasping points, and we select only plans maximizing the exposure goal, i.e. reaching a minimum visibility of 80%. Among these, we select the 3 plans minimizing F_{max} , which is the maximum recorded tissue force over the entire execution. Finally, we choose the optimal plan as the one with the lowest number of actions, thus minimizing interaction with tissue and obtaining the maximum economy of motion. We validate the full DEFRAS both in simulation (8 out of 17 pre-operative configurations) and on real setup (9 out of 17 pre-operative configurations).

A. Simulation Experiments

Each of the 8 pre-operative scenarios considered in our simulation experiments is characterized by a different configuration of APs (red outlines in Fig. 2) extracted as proposed in [37]. For each pre-operative scenario, we generate optimal plans following the strategy described above. Each pre-operative plan is executed by DEFRAS on two validation scenarios: the former with fewer APs (between 20-60%) and the latter with more APs (between 140-160%) with respect to the corresponding pre-operative configuration. These experiments emulate the realistic situation where initial knowledge of APs is not precise, thus the pre-operative plan might need adjustment due to mismatch between the current and the simulated environments. To run these experiments, we rely on two parallel simulations: the former (*pre-operative simulation*) represents the simulated environment in the scheme of Fig. 1, estimating the tissue forces and initialized with pre-operative APs; the latter (*actual simulation*) represents the validation scenario with either fewer or more APs than the pre-operative simulation, with simulated dVRK and cameras. Visibility in the actual simulation is computed by counting the number of visible ROI points from the simulated endoscope, assessed via ray-casting technique. The monitoring module continuously compares the tissue point clouds from the two simulations, triggering specific interventions according to Alg. 1. To assess the performance of the monitoring module, we compute the success rate of APs update with BANet as ratio between the number of times the update results in $\mu \leq \delta_1$ in Alg. 1 (no warning from monitoring) and the number of times BANet is triggered because $\mu > \delta_1$ (warning state).

We compare the performance of DEFRAS with respect to FRAS [16], the most similar framework in the surgical literature, in order to show the advantages of monitoring and learning modules in a deliberative architecture for autonomous robotic surgery. Therefore, we let FRAS perform autonomous TR on the actual simulations. It is worth recalling that FRAS does not rely on any pre-operative plan.

B. Real World Experiments

The performance of DEFRAS is evaluated on autonomous TR of silicone tissue with the real dVRK (Fig. 3). Similarly to [37], APs are defined by stitching the tissue on an attachment board with regularly spaced holes every 10 mm. The point cloud of the tissue in real environment is acquired after color-based segmentation and decimation of the raw

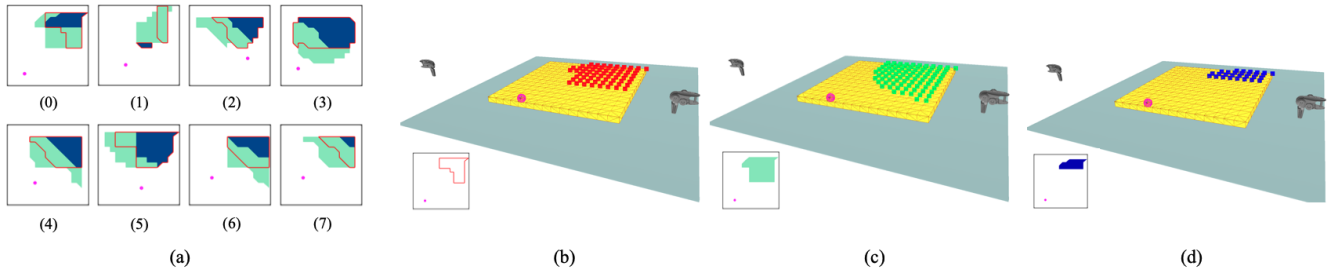


Fig. 2. (a) Configurations of APs used for the experiments in simulation. The red outline delimits the pre-operative configuration used to generate the plan; considered actual simulations are in blue (\ominus) and green (\oplus); magenta sphere represents the ROI. (b) Zoom in on the pre-operative configuration for case 0; (c) Zoom in on the \oplus actual configuration for case 0; (d) Zoom in on the \ominus configuration for case 0.

point cloud from Realsense as in [34], and compared with the point cloud from the simulated environment for monitoring. Visibility in the real environment is computed detecting the ROI (with known size and position) from the endoscopic camera images with color-based segmentation and considering the percentage of visible pixels, as in [25].

We consider three real world scenarios (A, B and C), obtained by stitching the tissue to the attachment board with 3 different configurations of APs (light blue in Fig. 3b). For each scenario, we consider 3 pre-operative models, each characterized by a different pre-operative knowledge of APs: (1) fewer APs (between 20-60%), (2) more APs (between 140-160%), and (3) same APs than/as the real scenario (Fig. 3b). For each of them, we generate the pre-operative plan in simulation, and test DEFRAS on the corresponding real configuration with the dVRK.

Moreover, we analyze the performance of our framework with respect to two state-of-the-art implementations of autonomous TR on real setups. In particular, we select the work by Pore et al. [25], which proposes reinforcement learning to generate an optimal motion-level policy for retraction, and the one by Attanasio et al. [29], which relies on image-guided control to reveal a ROI behind tissue flap. Similarly to us, both these approaches evaluate task success as percentage of ROI exposure upon task completion.

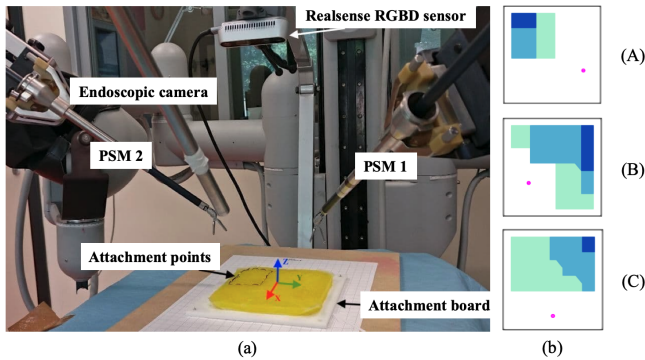


Fig. 3. (a) The setup for TR with dVRK. (b) Configurations of APs used for the real world experiments. The light blue region in each subimage represents the configuration of APs in the real setup (\ominus). Blue and green regions indicate pre-operative configurations with fewer APs (\ominus) and more APs (\oplus) than the real configuration. Magenta sphere represents the ROI.

VI. RESULTS AND DISCUSSION

A. Simulation Experiments

The optimal plans for each pre-operative configuration successfully expose the ROI with a maximum of 6 actions, with the force exceeding the safety threshold ϵ at most once, but always keeping F_{max} below 0.4 N.

Table I details the results obtained with DEFRAS and FRAS on scenarios with fewer and more APs than the corresponding pre-operative configuration. With DEFRAS, the ROI is successfully exposed in 13 out of 16 cases, even when the pre-operative information (i.e., knowledge of APs) is imprecise. Results show that, even when starting from pre-operative plans which minimize the tissue force, the force threshold can be exceeded during task execution, hence force feedback is needed to prevent dangerous tissue solicitation. The tissue force reaches higher values (> 1 N) mostly in scenarios with more APs than the pre-operative configuration (\oplus in Table). This is motivated by the fact that optimal plans might select a grasping point close to APs that are unknown pre-operatively. In this case, monitoring detects a discrepancy between the pre-operative simulation and the actual simulation, and triggers learning for APs update close to the current grasping point, resulting in high force estimate. The results of our framework also show that when pre-operative APs overestimate actual APs (\ominus in Table), fewer actions are needed with respect to the case when APs are underestimated. Therefore, overestimating APs in the pre-operative phase leads to more optimized task executions, reducing (potentially unsafe) interaction with the tissue.

From the comparison with FRAS, the presence of the pre-operative plan in DEFRAS proves fundamental to optimize executions in terms of tissue solicitation. Monitoring and learning modules in DEFRAS also improves performance. Whenever the mismatch between the pre-operative simulation and the actual simulation causes the monitoring module to trigger update of APs, BANet is able to provide an updated estimate of APs that brings the monitoring metric below threshold, except cases 1 \oplus (task success, ROI is exposed) and 5 \ominus (the task fails). Overall, the integration of the proposed deliberative modules results in higher median visibility (100% vs. 93%), lower number of failed tasks (3 vs. 5) and less frequent overcoming of the force threshold (-25% times, with -30% median force) with respect to FRAS.

TABLE I

SIMULATION EXPERIMENTS WITH FEWER (\ominus) AND MORE (\oplus) APs THAN PRE-OPERATIVE CONFIGURATION. FOR EACH METRIC, UPPER ROW IS REFERRED TO DEFRAS (D), LOWER ROW TO FRAS (F). "BANet" ROW REPORTS THE SUCCESS RATE OF APs UPDATE WITH BANet IN DEFRAS.

		0		1		2		3		4		5		6		7		Median	IQR
		\ominus	\oplus	\ominus	\oplus	\ominus	\oplus	\ominus	\oplus	\ominus	\oplus	\ominus	\oplus	\ominus	\oplus	\ominus	\oplus		
Visibility [%]	D	100	100	100	100	100	0	9	95	100	100	0	100	100	100	100	100	100	98-100
	F	100	90	100	0	0	0	90	0	100	100	100	0	97	93	93	93	93	0-100
Actions	D	4	11	6	6	6	30	38	12	5	26	29	25	5	30	5	5	8	5-26
	F	5	6	15	23	22	22	7	22	17	19	7	22	16	16	9	15	16	8-22
$F > \epsilon$	D	0	26	11	3	1	44	240	27	0	46	99	83	8	61	1	3	18	2-49
	F	0	0	29	33	32	28	11	40	21	29	0	28	27	18	20	12	24	11-29
F_{max} [N]	D	0.2	1.3	0.4	0.4	0.4	1.4	0.9	1.8	0.3	1.7	3.3	3.3	0.4	1.2	0.3	0.4	0.6	0.4-1.5
	F	0.3	0.3	1.0	1.0	1.1	1.0	0.4	1.2	0.9	0.9	0.3	1.1	0.9	0.9	1.2	0.9	0.9	0.8-1.0
BANet	D	1/1	1/1	-	0/1	2/2	1/1	6/6	-	1/1	2/2	11/12	-	1/1	4/4	2/2	-	-	-
	F																		

TABLE II

REAL WORLD EXPERIMENTS WITH DEFRAS. FOR EACH REAL WORLD CONFIGURATION (A, B, C), WE REPORT DETAILS OF TR EXECUTIONS WHEN UNDERESTIMATING REAL APs (\ominus), OVERESTIMATING REAL APs (\oplus) AND HAVING PERFECT KNOWLEDGE OF THE REAL APs (\ominus).

		A			B			C			Median	IQR
		\ominus	\oplus	\ominus	\ominus	\oplus	\ominus	\ominus	\oplus	\ominus		
Visibility [%]		99	95	100	100	100	71	99	83	96	99	95-100
Actions		13	15	9	24	7	8	14	7	7	9	7-14
$F > \epsilon$		11	20	8	0	8	20	8	8	6	8	8-11
F_{max} [N]		1.2	1.1	0.5	0.2	0.5	1.1	1.1	0.7	0.5	0.7	0.5-1.1
BANet		1/1	1/1	1/1	13/14	-	-	1/1	-	-	-	-

We notice that FRAS is able to solve the task when DEFRAS fails in configurations $3\ominus$ and $5\ominus$. However, in both cases the monitoring module triggers APs update and the force threshold is overcome more times than all other configurations. This suggests that the pre-operative knowledge of APs is particularly far from reality in these cases, meaning that the estimate of the force from the pre-operative simulation provided by FRAS might be unreliable.

B. Real World Experiments

Table II presents the results of experiments on the real setup. Obtained results confirm that whenever pre-operative APs underestimate real APs, the monitoring and learning modules are triggered more frequently. The update of APs leads to successful task completion in all cases. Moreover, when the simulated and real APs match, the number of actions is usually minimum (except for configuration B, with 8 actions for \ominus and 7 for \oplus). This suggests that overall the optimal pre-operative plan leads to better performances when APs are perfectly known.

In the video attachment, we show the executions in configuration A. The video shows that DEFRAS is able to re-plan and complete the task in case of grasping close to unknown APs (\ominus) and null visibility after optimal plan (\oplus). Re-planning proves able to successfully expose the ROI in both cases. Monitoring module detects a mismatch also in \ominus case, due to uncertainty introduced by real cameras and the calibration procedure. However, DEFRAS is able to cope with such situation and successfully complete the task.

The average visibility achieved by DEFRAS is 94%, which is aligned with the 90% visibility obtained by [25], and

improves the one reported in [29] (83%). Both reference works do not implement monitoring and learning.

VII. CONCLUSIONS

In this paper we proposed DEFRAS, a deliberative framework for autonomous robotic surgery. We introduced the fundamental capabilities of monitoring and learning in the context of a logic-based framework for autonomous interpretable planning, to deal with the uncertainty of anatomical environments. Experiments on a representative soft tissue retraction task have shown that our framework can learn and refine available pre-operative knowledge about APs in an online simulation, relying on constant monitoring of sensors and simulation data. In this way, any pre-operative plan can be adapted, and tissue force information from the simulation remains reliable for task completion. Compared with a framework lacking deliberative functions and recent tissue retraction implementations, results show that the percentage of task success increases, especially when prior APs are overestimated. Moreover, reliance on a pre-operative patient-specific plan brings benefits in terms of minimal manipulation of the tissue, hence reduced risk of tissue damage.

The modularity of DEFRAS will allow us to improve individual deliberative capabilities in order to deal with more realistic and challenging surgical scenarios in future works. Motion adaptation and learning as in [25], [40] will be implemented to replicate expert surgeons' dexterity. Furthermore, we plan to improve task knowledge with pre-operative [41] and intra-operative learning [42], [43]. Finally, we will cope with surgical-specific problems as vision occlusion.

REFERENCES

- [1] G.-Z. Yang *et al.*, “Medical robotics—regulatory, ethical, and legal considerations for increasing levels of autonomy,” *Science Robotics*, vol. 2, no. 4, p. 8638, 2017.
- [2] C. D’Ettorre *et al.*, “Accelerating surgical robotics research: A review of 10 years with the da vinci research kit,” *IEEE Robotics & Automation Magazine*, 2021.
- [3] A. Attanasio *et al.*, “Autonomy in surgical robotics,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 651–679, 2021.
- [4] F. Ingrand and M. Ghallab, “Deliberation for autonomous robots: A survey,” *Artificial Intelligence*, vol. 247, pp. 10–44, 2017.
- [5] M.-C. Fiazza and P. Fiorini, “Design for interpretability: Meeting the certification challenge for surgical robots,” in *2021 Int. Conf. Intelligence and Safety for Robotics (ISR)*. IEEE, 2021.
- [6] P. Kazanzides *et al.*, “An open-source research kit for the da vinci® surgical system,” in *2014 IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2014, pp. 6434–6439.
- [7] T. D. Nagy and T. Haidegger, “A dvrk-based framework for surgical subtask automation,” *Acta Polytechnica Hungarica*, pp. 61–78, 2019.
- [8] S. Sen *et al.*, “Automating multi-throw multilateral surgical suturing with a mechanical needle guide and sequential convex optimization,” in *2016 IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4178–4185.
- [9] A. Shademan *et al.*, “Supervised autonomous robotic soft tissue surgery,” *Science translational medicine*, vol. 8, no. 337, pp. 337ra64–337ra64, 2016.
- [10] J. Y. Wu, P. Kazanzides, and M. Unberath, “Leveraging vision and kinematics data to improve realism of biomechanical soft tissue simulation for robotic surgery,” *Int. Journ. Computer Assisted Radiology and Surgery*, vol. 15, no. 5, pp. 811–818, 2020.
- [11] F. Liu *et al.*, “Real-to-sim registration of deformable soft tissue with position-based dynamics for surgical robot autonomy,” in *2021 IEEE Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 12 328–12 334.
- [12] G. De Rossi *et al.*, “Cognitive robotic architecture for semi-autonomous execution of manipulation tasks in a surgical environment,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 7827–7833.
- [13] P. Berthet-Rayne *et al.*, “Hubot: A three state human-robot collaborative framework for bimanual surgical tasks based on learned models,” *2016 IEEE Int. Conf. Robotics and Automation (ICRA)*, pp. 715–722, 2016.
- [14] A. Murali *et al.*, “Learning by observation for surgical subtasks: Multilateral cutting of 3d viscoelastic and 2d orthotropic tissue phantoms,” in *2015 IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1202–1209.
- [15] B. Kehoe *et al.*, “Autonomous multilateral debridement with the raven surgical robot,” in *2014 IEEE Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1432–1439.
- [16] D. Meli *et al.*, “Autonomous tissue retraction with a biomechanically informed logic based framework,” in *2021 International Symposium on Medical Robotics (ISMR)*. IEEE, 2021, pp. 1–7.
- [17] M. Bratman, *Intention, Plans, and Practical Reason*. Cambridge: Cambridge, MA: Harvard University Press, 1987.
- [18] M. P. Georgeff and A. L. Lansky, “Reactive reasoning and planning,” in *AAAI*, vol. 87, 1987, pp. 677–682.
- [19] M. E. Bratman *et al.*, “Plans and resource-bounded practical reasoning,” *Computational intelligence*, vol. 4, no. 3, pp. 349–355, 1988.
- [20] E. Bolisani and C. Bratianu, “Knowledge strategy planning: an integrated approach to manage uncertainty, turbulence, and dynamics,” *Journal of Knowledge Management*, 2017.
- [21] F. Lally and P. Jannin, “Surgical process modeling: a review,” *Int. Journ. Computer Assisted Radiology and Surgery*, vol. 9, no. 3, pp. 495–511, 2014.
- [22] V. Lifschitz, “Answer set programming and plan generation,” *Artificial Intelligence*, vol. 138, no. 1-2, pp. 39–54, 2002.
- [23] F. Calimeri, W. Faber, M. Gebser, G. Ianni, R. Kaminski, T. Krennwallner, N. Leone, M. Maratea, F. Ricca, and T. Schaub, “Asp-core-2 input language format,” *Theory and Practice of Logic Programming*, vol. 20, no. 2, pp. 294–309, 2020.
- [24] M. Gebser, R. Kaminski, B. Kaufmann, M. Ostrowski, T. Schaub, and S. Thiele, “A user’s guide to gringo, clasp, clingo, and iclingo,” 2008.
- [25] A. Pore, E. Tagliabue, M. Piccinelli, D. Dall’Alba, A. Casals, and P. Fiorini, “Learning from demonstrations for autonomous soft-tissue retraction,” in *2021 International Symposium on Medical Robotics (ISMR)*. IEEE, 2021, pp. 1–7.
- [26] A. Pore *et al.*, “Safe reinforcement learning using formal verification for tissue retraction in autonomous robotic-assisted surgery,” in *2021 IEEE/RSJ Int. Conf. Intell. Robots and Systems (IROS)*. IEEE, 2021, pp. 0–6.
- [27] S. Patil and R. Alterovitz, “Toward automated tissue retraction in robot-assisted surgery,” in *2010 IEEE Int. Conf. Robotics and Automation*. IEEE, 2010, pp. 2088–2094.
- [28] R. Jansen *et al.*, “Surgical retraction of non-uniform deformable layers of tissue: 2d robot grasping and path planning,” in *2009 IEEE/RSJ Int. Conf. Intell. Robots and Systems*. IEEE, 2009, pp. 4092–4097.
- [29] A. Attanasio *et al.*, “Autonomous tissue retraction in robotic assisted minimally invasive surgery—a feasibility study,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6528–6535, 2020.
- [30] E. Tagliabue *et al.*, “Soft tissue simulation environment to learn manipulation tasks in autonomous robotic surgery,” in *2020 IEEE Int. Conf. Intell. Robots and Systems (IROS)*. IEEE, 2020.
- [31] C. Shin *et al.*, “Autonomous tissue manipulation via surgical robot using learning based model predictive control,” in *2019 Int. Conf. Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3875–3881.
- [32] S. A. Pedram *et al.*, “Toward synergic learning for autonomous manipulation of deformable tissues via surgical robots: An approximate q-learning approach,” in *2020 8th IEEE RAS/EMBS Int. Conf. for Biomedical Robotics and Biomechanics (BioRob)*. IEEE, 2019, pp. 878–884.
- [33] S. Nikolaev and S. Cotin, “Estimation of boundary conditions for patient-specific liver simulation during augmented surgery,” *Int. Journ. Computer Assisted Radiology and Surgery*, 2020.
- [34] E. Tagliabue *et al.*, “Intra-operative update of boundary conditions for patient-specific surgical simulation,” in *Int. Conf. Medical Image Computing and Computer Assisted Intervention*. Springer International Publishing, 2021.
- [35] J. Wang *et al.*, “Influence of clamping stress and duration on the trauma of liver tissue during surgery operation,” *Clinical Biomechanics*, vol. 43, pp. 58–66, 2017.
- [36] M. Machin *et al.*, “Smof: A safety monitoring framework for autonomous systems,” *IEEE Trans. on Systems, Man, and Cybernetics: Systems*, vol. 48, no. 5, pp. 702–715, 2016.
- [37] E. Tagliabue *et al.*, “Data-driven intra-operative estimation of anatomical attachments for autonomous tissue dissection,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1856–1863, 2021.
- [38] A. Roberti *et al.*, “Improving rigid 3-d calibration for robotic surgery,” *IEEE Trans. on Medical Robotics and Bionics*, vol. 2, no. 4, pp. 569–573, 2020.
- [39] F. Faure *et al.*, “Sofa: A multi-model framework for interactive physical simulation,” in *Soft tissue biomechanical modeling for computer assisted surgery*. Springer, 2012, pp. 283–321.
- [40] M. Ginesi *et al.*, “Overcoming some drawbacks of dynamic movement primitives,” *Robotics and Autonomous Systems*, vol. 144, p. 103844, 2021.
- [41] M. Bombieri *et al.*, “Automatic detection of procedural knowledge in robotic-assisted surgical texts,” *Int. Journ. Computer Assisted Radiology and Surgery*, pp. 1–9, 2021.
- [42] D. Meli *et al.*, “Inductive learning of answer set programs for autonomous surgical task planning,” *Machine Learning*, pp. 1–25, 2021.
- [43] D. Meli *et al.*, “Towards inductive learning of surgical task knowledge: A preliminary case study of the peg transfer task,” *Procedia Computer Science*, vol. 176, pp. 440–449, 2020.