Encountered-type haptic interface for virtual interaction with real objects based on implicit surface haptic rendering for remote palpation*

Alessandro Filippeschi[1], Filippo Brizzi[1], Emanuele Ruffaldi[1], Juan Manuel Jacinto[1], Carlo Alberto Avizzano[1]

***Abstract*— The shortage of physicians afflicting developed countries encourages engineers and doctors to collaborate towards the development of telemedicine. In particular, robotic systems have the potential for helping doctors making examination. A very common examination that can be the goal of a robotic system is palpation. Most of the robotics systems that have been developed for palpation present interesting features such as integrating augmented reality environments or allowing for hand free interaction. In this paper we present a novel palpation system that allows us to perform virtual palpation of real objects by means of a haptic and an augmented reality feedback. This system features an encountered-type haptic interface in which the haptic feedback is calculated by a collision detection algorithm that is based on online recording of the surface to be touched. The system allows the users to remove their hand from the haptic interface end-effector that follows the user's hand thanks to the tracking performed by a Leap Motion. We show that the system provides a natural interaction during the contact-non contact switch, a suitable force during indentation, and it allows to discriminate objects within the body through the haptic channel.**

I. INTRODUCTION

Today and in the following years, the increasing aging of the developed countries population requires and will require more and more medical examinations and interventions to be carried out [1]. This need causes a shortage of physicians that will become worse in the next years. In the recent years technology has proved and is proving to be a solution to fill the gap between the required amount of examinations and what physicians can actually do. In particular, robots showed to be valuable aids (e.g. see [2] for endoscopy, [3], [4]) for the doctors. Among the several applications of robot-aided examination, palpation has attracted the attention of both physicians and engineers in the last years.

Palpation is an examination in which the doctor's fingers and palms interact with the patient's abdomen. The doctors explores the organs and the other tissues beneath the skin checking whether any abnormality is present, such as organs bigger than they should be, nodules or other masses. Doctors adopt various techniques for abdominal palpation, varying the pushing location and the force.

Robotic systems were developed in the recent years to simulate palpation (e.g. [5], [6]). In [7], Inoue et al. propose a simple device for abdomen palpation that is composed of two sheets held by two wooden boards. The motion of the board as well as the tension in the two sheets allow for varying the stiffness perceived by the user. Although the system proved to reliably simulate the abdomen wall, it does not allow for the simulation of abnormalities under the skin surface. The HIRO system [8] features palpation of deformable tissue (breast palpation) based on a finite element model that allows to provide haptic feedback by means of a robot. Five thimbles are attached to the fingers of the robotic hand to provide forces to the user's fingers. To the authors' knowledge, the system does not currently feature visual feedback. Coles et al. propose in [9] a palpation and needle insertion system that integrates haptic feedback and virtual reality. The haptic feedback is provided by two Novint Falcon devices that actuate a tactile palpation end-effector. A camera placed on top of the users' hands track its motion. Finally, an augmented reality (AR) representation of the scene is co-located with the haptic feedback device, and it is displayed on a screen. The co-located haptic and visual feedback increase the fidelity of the system making it suitable for training. Currently, limitations of the system for abdomen palpation are the lack of representation of the real patient and the representation of abnormalities that are beneath the skin. The approach proposed by Diez et al. in [10] adopts also the encountered-type paradigm. They simulate the skin by means of a flat rubber sheet. On one side the user can touch the screen, whereas on the other side a robot provides haptic feedback according to a stiffness map that is calculated depending on an assumed tissue stiffness and on nodules' stiffness. The authors propose an algorithm that switches from position to impedance control depending on the interaction of the user with the screen. The user's hand is tracked by means of a marker-based optical tracking system. In [11] the combination of visual and haptic cues is further pushed, as multiple point haptic feedback is featured. The image of the patient is used for the visual feedback and for creating a domain in the space in which organs and abnormalities are immersed. The haptic feedback is then based on a blobby objects rendering approach. The visual display is based on a 2D image of the patient that is deformed according to the doctor's hand depth in the direction perpendicular to the image.

This paper presents the architecture and the control strategy of a novel system for virtual remote examination. Our approach combines visual and haptics cues and it allows the doctors to move freely their hands in the space when they are out of the patient's body and to receive a force feedback only when they interact with the patient. An augmented reality (AR) environment facilitates natural interaction with the system. The AR environment is co-located with the real

*This work was supported by the European project ReMeDi, grant 610902)

[1]The authors are with the Percetual Robotics Laboratory, TECIP Institute, Scuola Superiore Sant'Anna, Pisa, Italy a.filippeschi@sssup.it

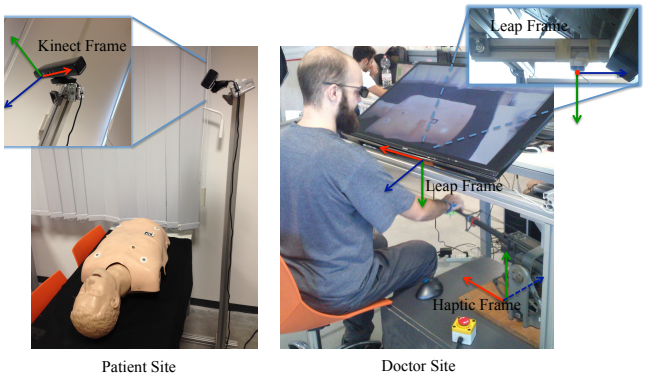

Fig. 1. The virtual palpation system. The components are shown along with the reference frames in which the gathered data are represented.

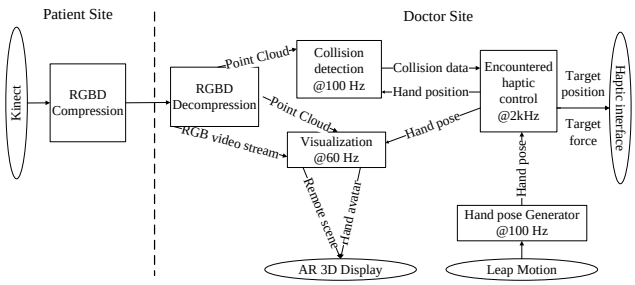

Fig. 2. The system architecture. The modules are shown along with their relationships.

hand and haptic interface and the doctor's hand is tracked regardless of the contact with the patient. The doctor is always displayed, in the AR environment, with an avatar of his/her hand that is suitably located with respect to the patient. Our system combines therefore the advantages of [9] (haptic rendering co-located with the AR environment) and of [10] (encountered-type interaction). Moreover, as for [11], in our approach the rendered force is the sum of two contributions: the force due to the indentation in the abdomen skin and the force due to internal structures. We progress from [11] approach as the first contribution is based not only on the doctor's current hand position, but also on an online scan of the patient's belly. Moreover, continuous tracking of the patient allows for a continuously updated 3D representation of the scene in the AR environment, in which the video stream of the patient and an avatar of the doctor's hand are displayed. This allows us to have a consistent haptic feedback even when the patient moves, thus handling real patients. In the future, this setting will allow us to extend our system to real remote examination, having a third (robotic) agent performing palpation guided by a doctor (see [12]). Therefore, this paradigm will be useful for training, for online (when combined with the robot) and offline remote examination.

After an overall presentation of the system, the components will be detailed. A particular focus is given to the control strategies that demonstrate the system flexibility that is needed to switch from a training setting to real remote examination. Then, a demonstrative experiment is presented to show the capabilities of the system.

II. METHOD

A. The palpation system

The virtual palpation system that we developed is shown in Figure 1. The system includes two locations: the patient's site (PS) and the doctor's site (DS). In the PS the patient lies on a table while he/she is tracked by a RGBD sensor (Microsoft Kinect® version 1). A computer manages the video and depth stream to the DS. In the DS, the doctor wears 3D glasses and seats in front of a frame that holds a 3D screen positioned horizontally. Under the screen the doctor's hand interacts with a high performance haptic interface (HI) [13] while his/her hand are tracked by means of a Leap Motion® mounted between the bottom of the screen and the HI. This solution has a sufficient tracking workspace and does not interfere with the contact of the hand with the haptic interface. In the DS a computer (PC-1) manages the video and depth stream coming from the PS. A second computer (PC-2) embeds the Matlab® XPC Target application that runs the control of the HI. Currently, the end-effector of the HI is a ball (6 cm diameter) that is used to handle the contact with either the fingers or the hand's palm. All the computers involved are Intel PC (Core i7 4770R 3.2 GHz, 8 GB RAM, embedded GPU) running Ubuntu Linux. The HI is a 3 DoFs robotic interface (see Figure 1). The workspace is a sphere sector that spans between 0.4 and 0.8 m from the center of rotation and include a barrel's rotation of $[-20,20]$ deg and $[-40,40]$ deg about the first two DoFs axes. The worst end-effector position resolution is 0.13 μm whereas the maximum continuous and peak force are 4N and 10N. Other HI could be used, given that they match the features of this one.

Thanks to the aforementioned components, the doctor interacts with a virtual model of the patient receiving a visual and a haptic feedback. The patient virtual model is based on a point cloud representation of the chest and abdomen surface obtained from the RGBD sensor. When the doctor's hand indents the skin the virtual force that is needed for the indentation is calculated. Moreover, this setup allows us to add components within the patient's virtual body such as organs and abnormalities, and to calculate the virtual force generated by interacting with these elements. The overall virtual force is then used to provide the doctor with a haptic feedback. When the doctor's hand is out of the body the HI end-effector is not in contact but still follows the doctor's hand. At the same time the doctor is provided with a 3D AR representation of the PS in which an avatar of the doctor's hand is superimposed to the PS scene.

B. System architecture

The system is made of several components acting either in the DS or the PS, they are represented along with their relationships in Figure 2. In the PS, the video and depth streams from the RGBD sensor are compressed and sent via

TCP to the DS. In DS PC-1 receives the information via TCP from the PS and runs the following modules:

- **visualization** (see section II-E), that provides the AR feedback based on the PS video stream and the doctor's hand position;
- **collision detection** (see section II-C), that exploits the point cloud from the PS and the doctor's hand position;
- **communication** with the patient site computer, the Leap Motion, and the XPC target computer for the haptic interface control.

The software framework used for the development of these modules, their latencies and the compression algorithms is called CoCo and is described in [14] and [15]. PC-2 receives data from PC-1 and runs the encountered haptic control that regulates both the haptic feedback and the doctor's hand position. Section II-D shows how the doctor's hand position is estimated depending on the contact with the HI.

Given the different computational burdens that the different tasks require, the aforementioned modules run at different rates. Figure 2 shows the rates the different modules run at along with the rate transitions that occur at the communication nodes. The constraints of each module will be described in the following sections. The relationships among reference frames that are needed for a consistent co-location of haptic and the visual feedback are obtained through the calibration procedure that are described in the following sections. Figure 2 reports the frames of each hardware/software component. In particular, the calibration procedure provides the transformation between the Kinect (K) and the Leap Motion (L) frames, and between the L and the haptic interface (H) frames.

C. Collision detection

The collision detection module is based on an implicit surface generated from the point-cloud using a KD-tree. This is a spatial structure well suited for noisy and varying point clouds, while other solutions can be employed for more rigid geometries [16]. Two parallel components realize this module: the first component receives the point cloud representation of the scene and uses it to build a KD-Tree. The KD-Tree is then passed to the second component which is in charge of providing the collision detection information. This component also receives the position of the hand, estimated by the HI end-effector to limit the sources of noise at the Kinect information only, converts it in the K frame, and uses it to query the KD-Tree. A radius search is performed on the point cloud gathering up to N points in a radius of R meters from the hand position. The system than performs a radius search, with the same parameters, on each of the previous points and uses the results to find the normal of each of the firsts points. The component uses the new information to calculate the implicit surface identified by the N points executing the algorithm in [17] and using as weight function the soft-objects function $C(r)$ described by Wyvill in [18]. Once the implicit function and the normal are known the component uses the algorithm described in [19] to calculate the position of the proxy in the scene. The proxy position is calculated at 100 Hz, converted in L frame and passed via UDP at the PC-2 together with the distance from the surface and a flag to notify the system whether the end-effector is over or below the implicit surface.

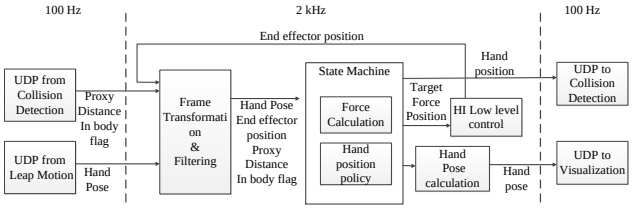

Fig. 3. The encountered haptic module along with its main components.

D. Encountered haptic control

This module (see Figure 3) manages the hand contact policy, determines the doctor's hand position and provides the doctor with a haptic feedback. It is implemented as a Matlab Simulink model that runs in external modality in the PC-2 XPC target at 2kHz frequency.

In the following ${}^{A}\mathbf{pB}_{c}$ means that the position $\mathbf{p}$ has been obtained by the device $\mathbf{B}$, is associated to the object c and is written in the frame A. If a variable has been inferred it has no capitol letter in its name. Each frame axis direction is associated to the unit vectors $\mathbf{i}$, $\mathbf{j}$, $\mathbf{k}$ for the x, y and z axes respectively. The variables available for the module are:

- ${}^{L}\mathbf{pL}_{h}$ and ${}^{L}TL_{h}$ i.e. the hand position and pose homogeneous matrix obtained from the Leap motion.
- ${}^{L}\mathbf{p}_{p}$ i.e. the proxy position from the collision detection.
- i_b i.e. a Boolean variable that is true when the doctor's hand is within the body
- ${}^{H}\mathbf{pH}_{E}$ i.e. the end-effector position provided by the HI.

The first three variables are available at 100 Hz and they are received via UDP. These variables are firstly converted to 2kHz frequency and then transformed in the HI frame by means of a calibration procedure that is described in the following to obtain ${}^{H}\mathbf{pL}_{h}$ and ${}^{H}\mathbf{p}_{p}$. ${}^{H}\mathbf{pH}_{E}$ is available from the HI sensors at 2kHz. Being based on ${}^{L}\mathbf{pH}_{h}$, i_b showed to be stable enough to be used without filtering. A Kalman filter was applied to ${}^{L}\mathbf{pL}_{h}$ and ${}^{L}\mathbf{p}_{p}$ to have a smoother signal. Given the low update rate of these variables, we exploited the previous estimates of these variables to predict the variable evolution in the 20 time steps between two updates.

The module returns the doctor hand position ${}^{L}\mathbf{pH}_{h}$ that is used for the collision detection and the doctor's hand pose ${}^{L}TH_{h}$ that is used by the visualization module.

$$ {}^{L}TH_{h} = \begin{bmatrix} {}^{L}RL_{h} & {}^{L}\mathbf{pH}_{h} \\ \mathbf{0} & 1 \end{bmatrix} \tag{1} $$

Moreover it embeds the low level control of the HI that features position and force control given a desired position of the end-effector and a desired force $\mathbf{F}$ to be displayed at the end-effector. The force control is an open loop control that provides the currents the motor must be supplied with.

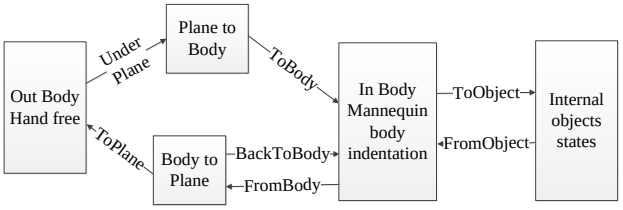

Fig. 4. The state machine that regulates the haptic feedback and the doctor's hand position.

The position control is a proportional-derivative controller based on ${}^{H}\mathbf{pH}_{E}$.

The calibration procedure consists of putting the hand on the end-effector where we desire to have the contact with the skin. ${}^{L}\mathbf{pL}_{h}$ and ${}^{H}\mathbf{pH}_{E}$ are recorded in this configuration to obtain ${}^{L}\mathbf{pL}_{h}^{0}$ and ${}^{H}\mathbf{pH}_{E}^{0}$ that are used as offsets for the following position transformations

- Leap to Haptic frame (L2H)

$$ {}^{H}\mathbf{p} = R_{L}^{H}({}^{L}\mathbf{p} - {}^{L}\mathbf{pL}_{h}^{0}) + {}^{H}\mathbf{pH}_{E}^{0} \tag{2} $$

- Haptic to Leap frame (H2L)

$$ {}^{L}\mathbf{p} = R_{H}^{L}({}^{H}\mathbf{p} - {}^{H}\mathbf{pH}_{E}) + {}^{L}\mathbf{pL}_{h}^{0} \tag{3} $$

where

$$ R_{L}^{H} = R_{H}^{L} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix} \tag{4} $$

We managed the transition between being out of the body and within the body by means of a finite state machine that is shown in Figure 4. Each state of the machine is labeled as s_f. In the *Out Body* condition the HI is in position control. The target hand position is

$$ {}^{H}\mathbf{p}_{t} = {}^{H}\mathbf{p}_{p} + [0\,\delta_{o}\,0]^{T} + \Delta^{H}\mathbf{pH}_{E} \tag{5} $$

where $\delta_o = 1\,\text{mm}$ is a small vertical offset that guarantees undesired transitions to the *In body*. $\Delta^{H}\mathbf{pH}_{E} = {}^{H}\mathbf{pH}_{E} - {}^{H}\mathbf{p}_{p}$ is calculated during the last sample before the transition from the *In body* state and allows us to avoid discontinuities in the position control of the HI just after the transition. The target force $\mathbf{F}$ is set to zero during this phase. Finally, in this state ${}^{L}TH_{h} = {}^{L}TL_{h}$. In the *In body* state the HI is in force control. This is obtained by setting ${}^{H}\mathbf{p}_{t} = {}^{H}\mathbf{pH}_{E}$. We define ${}^{H}\mathbf{p}_{p}^{0}$ as the proxy position just after the system enters the *In body* state. Hence ${}^{H}\mathbf{p}_{p}^{0}$ is the hand's first contact point on the undeformed skin. The target force $\mathbf{F}$ is then calculated depending on the indentation vector $\mathbf{d}$ with respect to the first contact point. This condition represents well the fact that during palpation the perceived force depends on how the skin has been stretched and indented with respect to the initial contact point (the hand does not slide on the skin)

$$ \mathbf{F} = [k_{x}(\mathbf{d}_{x})k_{y}(\mathbf{d}_{y})k_{z}(\mathbf{d}_{z})]^{T}\mathbf{d} \tag{6} $$

where

$$ k_{i}(d_{i}) = \alpha(d_{i}^{2} + d_{i}) \quad \text{i = x,y,z} \tag{7} $$

and

$$ \mathbf{d} = {}^{H}\mathbf{pH}_{E} - {}^{H}\mathbf{p}_{p}^{0} \tag{8} $$

This approach can be easily extended by continuously updating the proxy as in the algorithm presented in [19] where $\mathbf{d} = {}^{H}\mathbf{pH}_{E} - {}^{H}\mathbf{p}_{p}$ and k_i calculated according to 7. In case a collision occurs with other objects within the body, the force due to the collision can be calculated according to one of the previous two modalities and it is managed by means of further states. The overall target force is then calculated as the sum of the force due to the indentation $\mathbf{F}$ and the collision force with each of the objects. Currently a sphere and a cylinder are available to include objects within the body. Since the transition between the *In body* state and the *Out Body* state may cause discontinuities in the hand representation we sum the difference between the ${}^{L}\mathbf{pH}_{h}$ estimations occurring at the state transition. ${}^{L}RH_{h}$ is extracted from ${}^{L}TL_{h}$. Whereas the transitions among states within the body are simply dependent on the hand position with respect to the inner objects, the transitions between the *In body* state and the *Out Body* state may cause discontinuities in the visualization, instability in the control of the HI and undesirable loops between the two aforementioned states. To prevent these effects, we set a timer to stayat least 0.3 s in each state and we added two intermediate states called *Body to Plane* and *Plane to Body* (see Figure 4). The first avoids to have loops between *In body* and *Out Body* and allows for a smooth leaving of *In body*. The *FromBody* transition occurs when the collision detection module sets the in-body flag i_b to zero and the distance from the surface d_s is $d_S < d_{thr}$. In the *Body to Plane* state there's no position control but a force is applied to the end-effector in order to go farther from the body. When the height of the plane out of the body is reached, the transition *ToPlane* occurs. In the *Out Body* state there is no force applied to the end-effector, and a position control forces the end-effector to move on a plane above the body according to ${}^{H}\mathbf{pL}_{h}$ x and z components. When the user's hand pushes the end-effector towards the body, moving it below the plane, the *Under Plane* transition takes place. Then, in the *Plane to Body* state there is neither position nor force control, and the end-effector waits to be pushed towards the body. Finally, we allow the users to go from the *Body to Plane* back to the *In Body* state in case they start pushing again towards the body before leaving the intermediate state.

E. Visualization

The visualization module displays on the 3D screen the remote scene as a 3D mesh, created from the point cloud provided by the Kinect sensor, augmenting it with a virtual hand model controlled by ${}^{L}TH_{h}$.

The mannequin and therefore its 3D representation don't deform as the hand penetrates their surface causing the virtual hand to disappear beneath the mannequin mesh as soon as the indentation exceeds few centimeters. To avoid this inconvenience the position of the hand is moved, during the palpation, according to the position of the proxy ${}^{L}\mathbf{p}_{p}$, while the orientation is still ${}^{L}RL_{h}$. We set a zero-hold policy when ${}^{L}\mathbf{p}_{p}$ is not tracked.

To improve the alignment between the visual and the haptic feedback the texture of the virtual hand gradually shifts its color towards a red shade the deeper the end-effector is inside the surface.

III. TEST AND RESULTS

A. Assessment tests

Two healthy volunteers tested the system for a preliminary assessment of its usability. They carried out an experiment aimed at checking that forces are correctly displayed, thus allowing the participant to naturally interact with the body and to identify structures within the body. These goals are necessary steps to allow the final user to perform a correct diagnosis. The protocol is composed of four trials, in the first they had to indent the patient's skin in different points, thus verifying whether the system provides a natural interaction when switching between contact and non-contact with the patient. In the second trial, the volunteers had to enter the patient's body in a specific point and try to move within the patient's body. This trial allows us to verify whether equation 6 provides a suitable feedback for palpation. In the third trial, the volunteers were asked to interact with a cylinder (radius $r_c = 0.03$ m, height $h_c = 0.05$ m) lying 0.03 m under the skin. The cylinder's stiffness was set to 800 N/m in order to facilitate the volunteers identifying it. In the fourth trial they were left free to interact with the virtual patient looking for abnormalities within the body. In all the four trials we used the second modality (see equation 6) setting $\alpha = 1500$ (that corresponds to $k_x = k_y = k_z = 160$ N/m and $\|\mathbf{F}_b\| = 9.5$ N when the penetration is 0.06 m. In all the trials the collision detection module used $R = 0.1$ m and $N = 10$ for the radius search of the KD-tree, whose leaves include at most 10 points, and an influence radius of 0.1 m in the implicit surface algorithm. We gathered a 300k point cloud from the Kinect and all the KD-tree leaves were checked.

B. Results and discussion

We report here the results that we obtained, a video that was submitted along with the paper, was recorded to better show the results. In the following figures we show only explanatory examples, but the conclusions are supported by all the data that we gathered. In the following figures the possible states s_f are: $s_f = 0$ i.e. *In body*; $s_f = 1$ i.e. *Body to Plane*; $s_f = 2$ i.e. *Out body*; $s_f = 3$ i.e. *Plane to Body*; $s_f = -3$ i.e. *Cylinder Upper Area*; $s_f = -2$ i.e. *Around the Cylinder*; $s_f = -1$ i.e. *Cylinder Lateral Area*. The first trial confirmed that the system manages the transition between *Out Body* and *In Body* conditions preserving a natural interaction. Figure 5 shows that there are no jumps in the position of the end-effector, meaning that the transition is fluid. When the user's hand is outside of the body the end-effector correctly tracks the hand's position being on a plane just over the body (see the stars in Figure 5). The position of the hand to be used in the AR environment, instead, moves in 3D according to ${}^{L}TL_h$.

The second trial shows that the force is correctly displayed to the user in all the directions during the indentation. Figure 6 shows the force evolution according to the indentation and the force transition when switching the condition of the interaction (i.e. between inside and outside of the body).

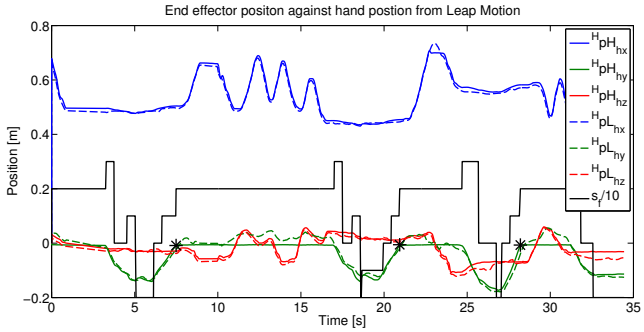

Fig. 5. First trial. Comparison of the end-effector position against the leap motion estimation of the hand position depending on the state s_f.

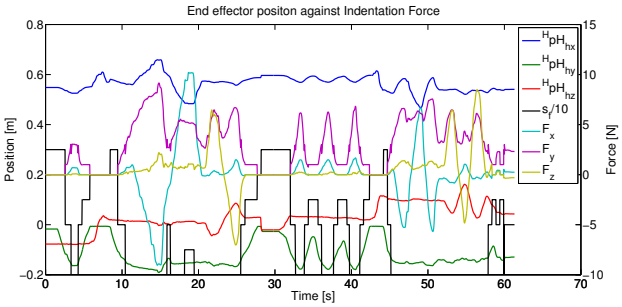

Fig. 6. Second trial. End-effector position against indentation force.

The third trial result shows that the volunteers were able to perceive the cylinder within the user body. Figure 7 shows the trajectories along with the forces that were perceived during the contact with the cylinder's upper and lateral area.

Figure 8 shows an example of interaction of the user with the environment when the goal was searching a cylinder. The system correctly performs in every s_f state, allowing the user to change the location of the exploration during the non-contact phases, to perceive the resistance of the abdomen tissue in the indentation phases and to perceive the force due the interaction with the cylinder.

As a final remark, the stiffness values were selected in these preliminary tests to ease the exploration of the belly. Increasing these values does not introduce any technical difficulty and more extensive tests will be carried out to make the system more specific for palpation. However, this first bench of tests demonstrate the usability and the capabilities of the system. We also recognize that a more reliable hand tracker could avoid the need for two intermediate states between the *In Body* and the *Out Body* states. We highlight, however, how the interaction with the system turns out to be very natural and fluid despite these states. The volunteers required indeed a short training to be able to use the system. We also mention here the nice trade off in the volunteers' strategy that we noted when searching for the cylinder in the fourth trial. The users firstly try to find the cylinder

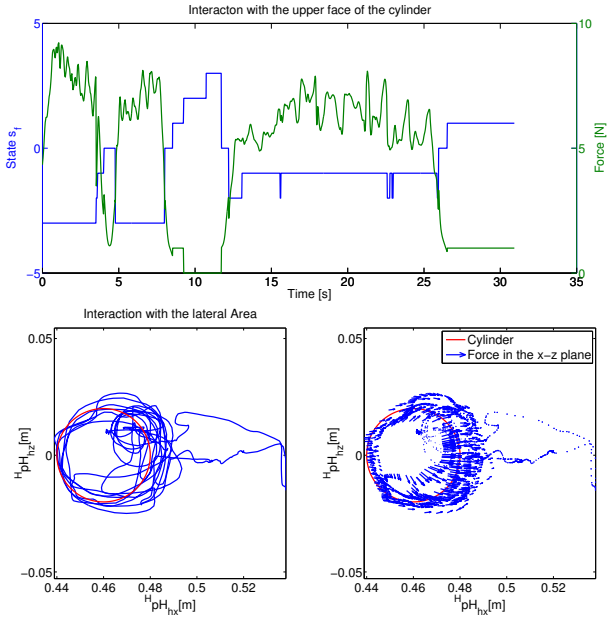

Fig. 7. Third trial. The upper figure shows the interaction with the upper area (see s_f). The lower figures shows the interaction with the lateral area.

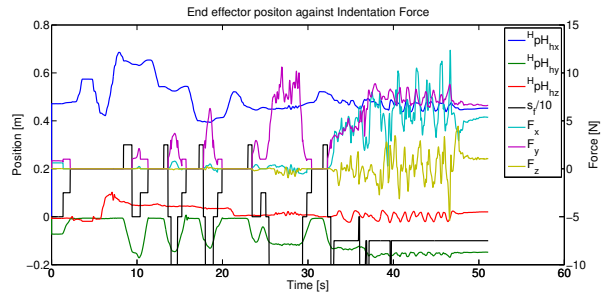

Fig. 8. Fourth trial. The search presents some of the elements of each of the previous trials: we note transitions form inside to outside with free hand and we not force exertion in different locations to find the cylinder.

from inside the body, thus experiencing increasing values of indentation force. When the indentation force was too high to keep searching, they exited from the body to start from another indentation point. This resembles what happens in palpation, in which the exploration is rather local around the first contact point with the skin. This aspect will be further investigated. Although there were not issues related to the resolution of the point cloud, we plan to test our system with the Kinect One sensor, in order to check possible improvements. In the present system Kinect was used because the newer sensor's drivers were not available for the Linux OS.

IV. CONCLUSION

This paper showed a novel system for virtual palpation of real objects. The system proved to manage correctly all of the phases of an encountered-type interaction and to provide the user with a force rendering that allows to perceive the indentation and to recognize abnormalities that are within the body. We will extend the validation setting and the system's parameters in order to be closer to actual palpation. Then we will assess the possibility for the users (including doctors) to recognize abnormalities with the body. This will enable the system to be used in a real telediagnosis setting, initially in a training environment, then with real patients.

REFERENCES

[1] R. M. Scheffler, J. X. Liu, Y. Kinfu, and M. R. Dal Poz, "Forecasting the global shortage of physicians: an economic-and needs-based approach," *Bulletin of the World Health Organization*, vol. 86, no. 7, pp. 516–523B, 2008.

[2] P. Valdastri, M. Simi, and R. J. Webster III, "Advanced technologies for gastrointestinal scopy," *Annual review of biomedical engineering*, vol. 14, pp. 397–429, 2012.

[3] P. Arbeille, A. Capri, J. Ayoub, *et al.*, "Use of a robotic arm to perform remote abdominal telesonography," *American Journal of Roentgenology*, vol. 188, no. 4, pp. W317–W322, 2007.

[4] C. Delgorge, Courrèges, *et al.*, "A tele-operated mobile ultrasound scanner using a light-weight robot," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 9, no. 1, pp. 50–58, 2005.

[5] R. L. Williams II, M. Srivastava, *et al.*, "The virtual haptic back for palpatory training," in *Prof of the 6th int. conference on Multimodal interfaces*. ACM, 2004, pp. 191–197.

[6] A. Talasaz and R. V. Patel, "Remote palpation to localize tumors in robot-assisted minimally invasive approach," in *IEEE ICRA*, 2012, pp. 3719–3724.

[7] K. Inoue, K. Ujiie, and S. Lee, "Development of haptic devices using flexible sheets for virtual training of abdominal palpation," *Advanced Robotics*, vol. 28, no. 20, pp. 1331–1341, 2014.

[8] T. Endo, H. Kawasaki, *et al.*, "Five-fingered haptic interface robot: Hiro iii," in *IEEE Worldhaptics*, 2009, pp. 458–463.

[9] T. R. Coles, N. John, D. A. Gould, and D. G. Caldwell, "Integrating haptics with augmented reality in a femoral palpation and needle insertion training simulation," *Haptics, IEEE Transactions on*, vol. 4, no. 3, pp. 199–209, 2011.

[10] S. P. Diez, E. B. Vander Poorten, G. Borghesan, and D. Reynaerts, "Towards palpation in virtual reality by an encountered-type haptic screen," in *Haptics: Neuroscience, Devices, Modeling, and Applications*. Springer, 2014, pp. 257–265.

[11] S. Yasmin and A. Sourin, "Virtual palpation for medical training in cyberworlds," in *Cyberworlds (CW), 2012 International Conference on*. IEEE, 2012, pp. 207–214.

[12] A. Peer, M. Buss, B. Stanczyk, *et al.*, "Towards a remote medical diagnostician for medical examination," 2014.

[13] M. Bergamasco, C. A. Avizzano, A. Frisoli, E. Ruffaldi, and S. Marcheschi, "Design and validation of a complete haptic system for manipulative tasks," *Advanced Robotics*, vol. 20, no. 3, pp. 367–389, 2006.

[14] E. Ruffaldi, A. Filippeschi, F. Brizzi, J. M. Jacinto, and C. A. Avizzano, "Encountered haptic augmented reality interface for remote examination," in *10th Symposium on 3D User Interfaces*. IEEE, 2015.

[15] L. Peppoloni, F. Brizzi, C. A. Avizzano, and E. Ruffaldi, "Immersive ros-integrated framework for robot teleoperation," in *10th Symposium on 3D User Interfaces*. IEEE, 2015.

[16] E. Ruffaldi, D. Morris, F. Barbagli, K. Salisbury, and M. Bergamasco, "Voxel-based haptic rendering using implicit sphere trees," in *Haptic Symposium*, March 2008, pp. 319–325.

[17] A. Adamson and M. Alexa, "Approximating and intersecting surfaces from points," in *Eurographics Symposium on Geometric Processing*, 2003, pp. 230–239.

[18] G. Wyvill, C. McPheeters, and B. Wyvill, "Data structure forsoft objects," *The visual computer*, vol. 2, no. 4, pp. 227–234, 1986.

[19] K. Salisbury and C. Tarr, "Haptic rendering of surfaces defined by implicit functions," in *ASME Dynamic Systems and Control Division*, vol. 61, 1997, pp. 61–67.