Image-Driven Stochastic Identification of Boundary Conditions for Predictive Simulation

Igor Peterlik[1], Nazim Haouchine[1], Lukáš Ručka[2], and Stéphane Cotin[1(✉)]

[1] Inria, Paris, France
`stephane.cotin@inria.fr`
[2] Faculty of Informatics, Masaryk University, Brno, Czech Republic

Abstract. In computer-aided interventions, biomechanical models reconstructed from the pre-operative data are used via augmented reality to facilitate the intra-operative navigation. The predictive power of such models highly depends on the knowledge of boundary conditions. However, in the context of patient-specific modeling, neither the pre-operative nor the intra-operative modalities provide a reliable information about the location and mechanical properties of the organ attachments. We present a novel image-driven method for fast identification of boundary conditions which are modelled as stochastic parameters. The method employs the reduced-order unscented Kalman filter to transform in real-time the probability distributions of the parameters, given observations extracted from intra-operative images. The method is evaluated using synthetic, phantom and real data acquired in vivo on a porcine liver. A quantitative assessment is presented and it is shown that the method significantly increases the predictive power of the biomechanical model.

Keywords: Boundary conditions · Stochasric data assimilation · Finite element method · Surgical augmented reality · Hepatic surgery

1 Introduction

The augmented reality (AR) has become an increasingly helpful tool for navigation in laparoscopic surgery and interventional radiology. For example, during the hepatic surgery, only a part of the liver surface is visible to the laparoscopic camera. In this case, the AR provides a visualization of the internal structures such as tumors and blood vessels in order to increase surgical accuracy and reduce the time of the intervention [1]. While the initial position of these structures is obtained from the pre-operative data, the actual configuration at any instant during the intervention requires a predictive model, which accounts for

the deformations and displacements that occur due to the surgical manipulations and other types of excitations. The existing works in this area [2,3] employ a patient-specific finite element (FE) biomechanical models reconstructed from the pre-operative data. During the intervention, the deformations computed by the models are driven by the image data provided by the laparoscopic camera.

Beside the geometry of the model, the boundary conditions (BCs) have a significant impact on the accuracy of the predictions computed by biomechanical models [4]. Nonetheless, the actual location and elastic properties of BCs, typically represented by ligaments and tendons, are also patient specific, but unlike the geometry of the organ, they cannot be identified directly from the pre-operative data. Moreover, the intra-operative modalities give only a partial and very inaccurate information about the BCs, which are often hidden from the intra-operative view.

Surprisingly, only few works address this issue. In [5], the authors propose a method estimating surface loads corresponding to BCs from the rest and deformed shape of an organ. As the complete surface must be available in both configurations, the method is not applicable without additional intra-operative scanning. In [6] a statistical atlas is used to transfer the positions of attachments to the actual geometry. In [7], the authors describe a database of liver deformations obtained from medical data, enriched with BCs using a simuation. However, due to the high variability in anatomy of some organs, the atlas-based methods provide only statistical information that is often difficult to instantiate for given patient.

In this paper, we propose a novel image-driven stochastic assimilation method to identify the BCs of a biomechanical model of a deformable body. The BCs are regarded as elastic springs attaching the organ to its surroundings. While we a priori select a surface region where the springs are located, we do not make any assumption about their elasticities: these are modelled as stochastic parameters with expected values initially set to zero. The reduced-order unscented Kalman filter is employed to transform the corresponding probability functions given the observations extracted from a video sequence. The method is evaluated using synthetic, phantom and real in-vivo data acquired by a laparoscopic camera. Beside qualitative assessment, a quantitative evaluation is given in terms of accuracy and stability as well as performance which is close to real time. To the best of our knowledge, this is the first method allowing for a fast identification of boundary conditions from image data.

2   Method

2.1   Overview of the Algorithm

The geometry of the biomechanical FE model is reconstructed from the pre-operative data, such as 3D CT. The FE mesh is initially registered to the first frame of the video sequence, using for example the method presented in [6]. Given the geometry of the model, two types of information must be available

in order to compute a deformation: the excitation (control input) and boundary conditions (BCs). In this paper, we consider the excitation to be known: it is represented by prescribed displacements, given for example by the actual position of a tool manipulating the deformable object. Since the control input is given by prescribed displacements and we suppose that the object has homogeneous material properties, the value of the material stiffness does not have a direct impact on computed deformations. In this paper, we employ the MJED formulation of the StVenant-Kirchoff hyperelastic material [8] which makes use of the full non-linear Green strain tensor while assuming the material linearity.

Whereas the control input is known, the boundary conditions are not given exactly: it is assumed that there is a surface region Σ represented by a set of mesh nodes which are attached to their rest positions by elastic springs. However, the elasticities of these springs, which can range from 0 (no attachment) to a high value (a stiff attachment), are not known a priori. Elasticity of each spring is treated as an independent stochastic parameter given by the Gaussian probability density function (PDF). Initially, $k_s \sim \mathcal{N}(0, \sigma)$ for each $s \in \Sigma$ with $\sigma > 0$ being the initial standard deviation. The proposed algorithm employs the reduced-order unscented Kalman filter (ROUKF) [9] to transform the PDF of each parameter k_s given the sequence of observations extracted from the video. The aim of the procedure is to minimize the standard deviation of each stochastic parameter k_s, thus performing the identification of boundary conditions (elastic attachments) needed to optimize the predictive power of the model.

2.2 Stochastic Identification of Boundary Conditions

Before giving the detailed description of the algorithm, we introduce the notation for entities used by the method. In what follows, a feature is a point extracted from the initial frame and tracked in all the following frames. Each feature f has two positions: p_f^t is its position given by the video frame at time step t, whereas q_f^t is the position of the feature attached to the model. The set $\mathcal{F}$ of all the features is divided into set of control features $\mathcal{C}$ and set of observation features $\mathcal{O}$, where $\mathcal{C} \cap \mathcal{O} = \emptyset$. The control features are typically (but not exclusively) selected close to the tool which manipulates the deformable body. During the assimilation, the control features are used to define the prescribed displacements of the model, while the observation features are used by the filter correction phase to compute the Kalman gain.

Further, Σ represents a region given by a set of S mesh nodes in which the stochastic springs are placed. Since we employ the simplex version of the ROUKF, there are $S + 1$ sigma points, i.e., in each prediction phase of the assimilation process, $S + 1$ evaluations of model are performed, one for each sigma point. The algorithm works as follows:

Initialization at time step $t = 0$**.** After initial registration of the model w.r.t. the first video frame, the extraction of control features and observation features is performed, resulting in positions p_c^0 and p_o^0. The corresponding model positions are set as $q_c^0 = p_c^0$ for each $c \in \mathcal{C}$ and $q_o^0 = p_o^0$ for each $o \in \mathcal{O}$ and

each of these positions is mapped to the model by computing the barycentric coordinates of each q_f w.r.t. the element of the model mesh in which the feature is located. The initial expected value of each stiffness parameters is set to 0 and the covariance matrix is set to a diagonal matrix $\sigma \mathbf{I}_{S \times S}$.

Prediction phase at time step $t > 0$. First, the image positions p_c^t of control features are obtained by tracking performed on the actual video frame. Second, $S + 1$ vectors of stiffness parameters $[k_{s_1}, \dots k_{s_S}]$ (sigma points) are sampled and $S + 1$ simulations are performed. Each simulation employs the absolute values of one sigma-point vector to parametrize the spring stiffnesses and the Newton-Raphson method is used to compute the quasi-static deformation of the object given the new positions $q_c^t = p_c^t$ of control features. Since the difference between q_c^{t-1} and q_c^t is typically small, the solution converges quickly. As the simulations performed for different sigma points are independent, they can be executed in parallel, thus accelerating the prediction phase. Finally, after the simulations for all sigma points are computed, the a priori expected value and covariance matrix are updated (see [9]).

Correction phase at time step $t > 0$. First, the image positions p_o^t of observation features are obtained from the actual video frame by a tracking algorithm. Second, the model positions q_o^t of observations features are updated using the model parametrized by the expected values of stiffness parameters compute by the prediction phase. The difference between p_o^t (the image-based positions) and q_o^t (the model-predicted positions) is used to compute the innovation, which is required to calculate the Kalman gain. Finally, this quantity is used to compute the a posteriori expected values and covariance matrix which together represent the transformed PDFs of spring stiffnesses.

3   Results

The proposed method is assessed using synthetic data, video capturing a controlled deformation of a silicon phantom and finally a video of deforming pig liver acquired by a laparoscopic monocular camera. As noted in Sect. 2.1, the value of object stiffness does not directly influence the deformations computed by the model. In all scenarios, we set Young's modulus to 5,000 kPa and Poisson ratio to 0.45 corresponding to a soft tissue. In the case of phantom and medical data, the features are extracted using FAST and tracked with the optical flow. The reported rates were obtained on a PC with CPU i7-6700 and 16 GB RAM.

For the sake of quantitative assessment performed for the synthetic and medical data, *assessment points* are selected inside or on the surface of the deformed object so that they coincide with neither the control nor the observation features. At every assimilation step $t > 0$, each assessment feature a has two different locations: p_a^t given by the image and q_a^t predicted by the model. At the end of the step, the distance error $\varepsilon_t = \max_a ||p_a^t - q_a^t||$ is calculated for each feature, quantifying the prediction error of the model.
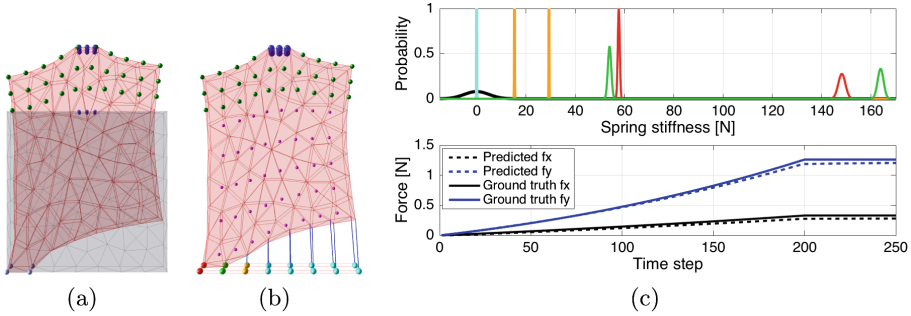
Fig. 1. Synthetic data: (a) rest shape (in grey) and deformation computed by the forward simulation with 4 fixed nodes. (b) Result of the stochastic assimilation showing the assessment points (magenta), control features (blue) and observation features (green). (c) top: Initial (black) and assimilated PDFs; the colors correspond to locations in subfigure (b). (c) bottom: Comparison of forces applied in the control features obtained by the forward simulation (solid) and stochastic simulation (dashed).

3.1   Evaluation Using the Synthetic and Phantom Data

In order to perform a quantitative validation of the method, we first created a model of 3D brick ($10 \times 10 \times 1$ cm) composed of 536 linear tetrahedra shown in Fig. 1a. The brick was deformed in a forward simulation: gradually increasing displacements were prescribed in four control features in order to emulate a tool pulling the brick upwards, while imposing the homogeneous Dirichlet boundary conditions on the left part of the base. The forward simulation was used to generate the observation positions p_o^t as well as the positions of assessment points p_a^t located inside the brick which thus serves as the ground truth.

Using the same control input and the observations generated by the forward simulation, stochastic data assimilation was performed, running at 35 FPS. In this case, the entire base of the brick was selected as the region Σ containing 16 nodes with associated stochastic springs. The transformation of the stochastic parameters is depicted in the upper part of Fig. 1c: initially, all the stochastic parameters $k_s \sim \mathcal{N}(0, 5)$ (the PDFs represented by black bell curve). The transformed PDFs given the synthetic observations are depicted in colors: comparison to the coloring of Σ nodes in Fig. 1b allows for matching between the location of the spring and its assimilated PDF: those with non-zero mean correspond to the springs located in the nodes constrained by the homogeneous Dirichlet BCs or in their neighbours. Therefore, despite the uniform initialization, the assimilation process correctly identifies the boundary conditions of the model. This is further confirmed by the error computed in the assessment points: during the assimilation process, the error ε_t computed for each t does not exceed 3.1 mm. This result was compared to a deterministic simulation where the homogeneous BCs are imposed either in all the nodes in the region Σ (resulting in $\max_t \varepsilon_t = 29$ mm) or alternatively no BCs are imposed at all ($\max_t \varepsilon_t = 35.7$ mm).

(a)   (b)
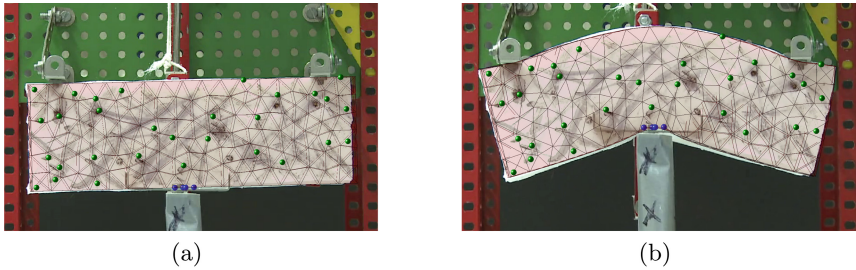
Fig. 2. Rest (a) and deformed (b) silicon phantom with superimposed model with boundary conditions assimilated by the stochastic method.

In the case of the synthetic data, it is possible to compute the forces applied in the control features for both the forward and stochastic simulations. The force plots depicted in the lower part of Fig. 1c show a very good agreement between ground truth and stochastic simulations. This result indicates that if the material properties of a deformable object are available, the stochastic framework allows for estimation of applied forces from the image data.

A qualitative evaluation of the method was performed using a silicon phantom deformed in a controlled environment as depicted in Fig. 2. The FE mesh of the silicon model was composed of 1216 tetrahedra ($16.3 \times 5.3 \times 1.3$ cm). The experimental images were acquired by a common digital camera. The top surface of the object was selected as the surface region Σ containing 53 mesh nodes and the goal of the assimilation was to identify the two obstacles located on the left and on the right side above the phantom. Visual comparison of Fig. 2a and b illustrates the accuracy of the assimilation process running at 10 FPS.

3.2   Evaluation Using In-vivo Data

The stochastic algorithm was evaluated using an intra-operative video sequence acquired in vivo on a porcine liver. The FE model of the liver was reconstructed from the pre-operative CT images acquired on the same pig, resulting in a mesh composed of 315 nodes. A short video sequence (about 7 s) was recorded using a monocular laparoscopic camera inserted into the porcine abdomen inflated by the gas. The video captures a deformation of one liver lobe manipulated with laparoscopic pincers. The region Σ containing 35 nodes was selected in the lower part of the organ which is occluded by other anatomical structures. The FE model was registered to the first frame of the video which was also used to automatically extract the features. These were projected on the surface of the registered model to obtain an initial alignment. Assessment points inside the liver would require an intra-operative scanner not available for this experiment. Therefore, we selected three surface features, each located in a different part of the liver lobe. All the features coloured according to the type are depicted in Fig. 3a. Figure 3b depicts the trajectory of each feature during the manipulation. This figure reveals that the motion of the liver lobe was not induced only by the
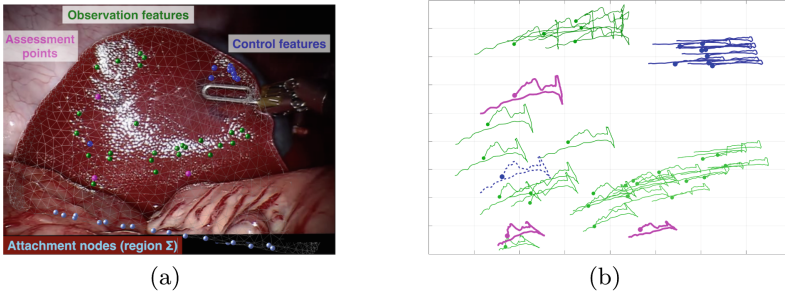
(a)   (b)

Fig. 3. (a) The first frame of the video sequence with features. (b) Trajectories of features during the manipulation. The control feature on the left side of the lobe (dashed trajectory) was employed only in the scenario 2.

tool: while the trajectories of the control points follow a horizontal line, the trajectories of both the observation and assessment features display important vertical perturbations probably due to the respiratory or cardiac motion. So we studied the behaviour of the assimilation in two different scenarios: beside using only the features in the right upper part of the lobe as control, in the second scenario, we added one more control feature located in the left part of the lobe.

For both scenarios, we compared the stochastic simulation with two deterministic simulations: in the first one, no boundary conditions were imposed during the entire manipulation, while in the second one, all the nodes in the region Σ were fixed to their rest positions. The error computed in each feature for the two scenarios and three types of simulations is plotted in Fig. 4. The evaluation shows that in both scenarios, the assimilation of boundary conditions results in improved predictive power of the model. In the first scenario, using no bound-
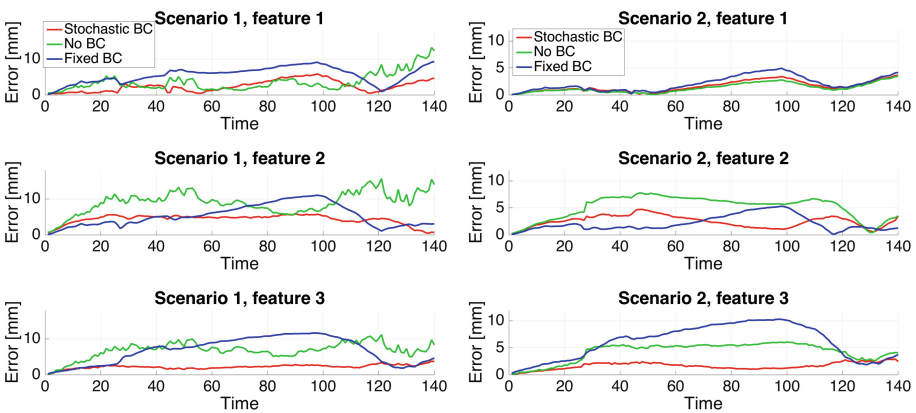

Fig. 4. Temporal evolution of ε_t computed for each t in 3 assessment points, plotted for two different scenarios: without (left) and with (right) additional control feature.

ary conditions leads to an unstable simulation, while fixing all the nodes lead to over-constrained simulation. Adding one more control feature significantly reduces the error in all types of simulations. Although in some time steps, the deterministic simulations yield better results, this behaviour is not consistent: e. g. fixing all the nodes would improve the error of feature 2 from $t = 30$ to $t = 70$, however, it would radically increase the error in feature 3. As for the performance, while the deterministic simulation runs at 40 FPS, the stochastic simulation with parallelized prediction runs at 15 FPS.

4   Discussion and Conclusion

In order to achieve an optimal and stable assimilation, it is necessary to *tune* the stochastic filter. Basically, the tuning is done by careful adjustment of the initial PDFs of the stochastic parameters. While for both synthetic and phantom data, we obtain stable and optimal assimilation by setting initially $k_s \sim \mathcal{N}(0, 5)$, much lower initial standard deviation must be chosen in the case of real data (0.0002 and 0.01 for scenario 1 and scenario 2, respectively). Beside the variance, the initial expected value of parameters can be initialized with some existing a priori knowledge. For example it can be initialized using a statistical anatomical atlas.

In the classical data assimilation and filtering theory, the observations utilized to correct the predictions are also treated as stochastic quantities. This approach accounts for the uncertainties related to the observation inputs. When considering the augmented reality, the observations are extracted from the intraoperative images typically suffering from high level of noise. Therefore, we believe that treating the observations as PDFs, would further improve the robustness of the stochastic simulation.

In this paper, we have presented a stochastic identification of boundary conditions for predictive simulation employed in the context of augmented reality during the hepatic surgery. The method was validated on different types of data, including a video capturing an in vivo manipulation of a porcine liver. Beside focusing on tuning the prediction and correction phases of the filter and more thorough validation of the algorithm, we plan to further improve the performance of the method by preconditioning and hardware acceleration.

References

1. Marescaux, J., Diana, M.: Next step in minimally invasive surgery: hybrid image-guided surgery. J. Pediatr. Surg. **50**(1), 30–36 (2015)
2. Haouchine, N., et al.: Image-guided simulation of heterogeneous tissue deformation for augmented reality during hepatic surgery. In: ISMAR 2013, pp. 199–208 (2013)
3. Suwelack, S., et al.: Physics-based shape matching for intraoperative image guidance. Med. Phys. **41**(11), 111901 (2014)
4. Bosman, J., Haouchine, N., Dequidt, J., Peterlik, I., Cotin, S., Duriez, C.: The role of ligaments: patient-specific or scenario-specific? In: Bello, F., Cotin, S. (eds.) ISBMS 2014. LNCS, vol. 8789, pp. 228–232. Springer, Cham (2014). doi:10.1007/978-3-319-12057-7_26

5. Peterlik, I., Courtecuisse, H., Duriez, C., Cotin, S.: Model-based identification of anatomical boundary conditions in living tissues. In: Stoyanov, D., Collins, D.L., Sakuma, I., Abolmaesumi, P., Jannin, P. (eds.) IPCAI 2014. LNCS, vol. 8498, pp. 196–205. Springer, Cham (2014). doi:10.1007/978-3-319-07521-1_21
6. Plantefève, R., et al.: Patient-specific biomechanical modeling for guidance during minimally-invasive hepatic surgery. Ann. Biomed. Eng. **44**(1), 139–153 (2016)
7. Johnsen, S.F., et al.: Database-based estimation of liver deformation under pneumoperitoneum for surgical image-guidance and simulation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9350, pp. 450–458. Springer, Cham (2015). doi:10.1007/978-3-319-24571-3_54
8. Marchesseau, S., Heimann, T., Chatelin, S., Willinger, R., Delingette, H.: Multiplicative jacobian energy decomposition method for fast porous visco-hyperelastic soft tissue model. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6361, pp. 235–242. Springer, Heidelberg (2010). doi:10.1007/978-3-642-15705-9_29
9. Moireau, P., Chapelle, D.: Reduced-order unscented Kalman filtering with application to parameter identification in large-dimensional systems. ESAIM: Control Optim. Calc. Var. **17**(2), 380–405 (2011)