



Editorial

Shaping the future through innovations: From medical imaging to precision medicine

Dorin Comaniciu^a, Klaus Engel^b, Bogdan Georgescu^a, Tommaso Mansi^{a,*}^a Medical Imaging Technologies, Siemens Healthcare Technology Center, Princeton, NJ, USA^b Medical Imaging Technologies, Siemens Healthcare Technology Center, Erlangen, Germany

ARTICLE INFO

Article history:

Received 5 April 2016

Revised 8 June 2016

Accepted 13 June 2016

Available online 15 June 2016

Keywords:

Photorealistic rendering

Artificial intelligence for image parsing

Heart valves modelling

Computational modelling

ABSTRACT

Medical images constitute a source of information essential for disease diagnosis, treatment and follow-up. In addition, due to its patient-specific nature, imaging information represents a critical component required for advancing precision medicine into clinical practice. This manuscript describes recently developed technologies for better handling of image information: photorealistic visualization of medical images with Cinematic Rendering, artificial agents for in-depth image understanding, support for minimally invasive procedures, and patient-specific computational models with enhanced predictive power. Throughout the manuscript we will analyze the capabilities of such technologies and extrapolate on their potential impact to advance the quality of medical care, while reducing its cost.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Medical imaging has impacted the practice of medicine during the recent decades, contributing to greatly improved disease diagnosis, treatment and follow-up. Image-guided, minimally invasive procedures are becoming more and more common in hospitals, replacing conventional surgery and allowing faster recoveries with fewer post-procedure complications. We anticipate that this trend will continue, medical imaging playing an increasingly important role towards moving precision medicine into clinical practice. By being able to characterize anatomy, physiology and metabolism of the patient, medical imaging enables precise, personalized procedures and predictive, patient-specific therapy selection and delivery.

In this paper we highlight a number of technologies that will most likely contribute to the success of medical imaging for the years to come, helping medical care to advance, while reducing its cost. In [Section 2](#) we discuss Cinematic Rendering, a 3D visualization technology that is capable of producing superb photorealistic images from traditional Computer Tomography (CT) or Magnetic Resonance (MR) volumes, thus potentially enhancing the conspicuity of pathologies. [Section 3](#) addresses the topic of next generation image understanding, which contributes to faster and more reproducible image reading, benefiting from the recent advances in machine learning and artificial intelligence. Furthermore, in

[Section 4](#) we discuss the real-time imaging needs in the operating room and focus on heart valve procedures, addressing both their planning and guidance. Finally, in [Section 5](#) we present patient-specific computational models that contribute to advances in diagnosis, patient stratification, therapy selection and therapy optimization. All images shown in the paper are images of real, living patients.

2. Cinematic rendering: photorealistic visualization of medical images

Efficient clinical decisions and procedures require the rapid appreciation of the relevant information contained within medical images. Even though medical image viewing based on multi-planar reconstruction (MPR) is still dominant in diagnostic imaging, the significance of three-dimensional visualization of medical data is rising. This is due to the fact that these methods allow much faster understanding of spatial anatomical structures and have the potential to increase the sensitivity and specificity of medical images. Especially medical professionals who are not trained in planar image viewing as well as patients benefit from such visualizations.

Recent advances in computer graphics have made interactive physically-based volume visualization techniques possible. Such techniques reproduce complex illumination effects in computer-generated images by mimicking the real-world interaction of light with matter. The results are physically plausible images that are often easier for the human brain to interpret, since the brain is trained to interpret the slightest shading cues to reconstruct shape and depth information. Such shading cues are often missing from

* Corresponding author.

E-mail address: tommaso.mansi@siemens.com (T. Mansi).



Fig. 1. Cinematic Rendering. Left: Original computed tomography (CT) data; Right: Cinematic rendering of the same dataset. Data courtesy of Israelitisches Krankenhaus, Hamburg, Germany.

computer generated images based on more simple geometric calculations such as ray casting.

We developed a physically-based volume rendering method called Cinematic Rendering (Engel 2016; Paladini 2015) which computes in real-time the interaction of visible photons with the scanned patient anatomy. The algorithm uses a Monte Carlo path tracing method to generate photorealistic or even hyper-realistic images by light transport simulation along hundreds or thousands of photons paths per pixel through the anatomy using a stochastic process (Fig. 1)

In traditional volume ray casting, only emission and absorption of radiant energy along a straight ray is considered. Radiant energy q_e is emitted at each point x' along the ray up to a maximum distance D .

$$L(x, \omega) = \int_0^D e^{-\tau(x, x')} q_e(x') dx' \quad (1)$$

The emitted radiant energy at each point is absorbed according to the Beer-Lambert law along the ray to the observer location with absorption coefficients σ_a .

$$\tau(x, x') = \int_x^{x'} \sigma_a(t) dt \quad (2)$$

Single scattering is usually modelled in traditional volume rendering using a surface shading model that considers local gradient information of the volume data (local illumination). While this integral can be easily solved numerically using a Riemann integral, the method neglects complex light paths with multiple scattering events and extinction of light (global illumination).

In contrast, the Monte Carlo path tracing integration method solves the following multi-dimensional and non-continuous rendering equation:

$$L(x, \omega) = \int_0^D e^{-\tau(x, x')} \sigma_s(x') \left[\int_{\Omega_{4\pi}} p(\omega, \omega') L_i(x', \omega') d\omega' \right] dx' \quad (3)$$

Eq. (3) determines the radiant flux (radiance) L at distance x received from the direction ω along a ray. We have to integrate the radiance scattered into that direction from all possible directions ω' at all points along the ray up to a maximum distance D . The optical properties of a relevant tissue are defined using the phase function $p(\omega, \omega')$, which describes the fraction of light travelling along a direction ω' being scattering into the direction ω . $L_i(x', \omega')$ is the radiance arriving a distance x' from direction ω' . In practice, we model scattering in different tissue types using a Henyey-Greenstein phase function and compute shading of implicit surfaces using a BRDF (bidirectional reflectance distribution function).

Radiance scattering into the direction ω is also absorbed and scattered out of the direction ω . This is modelled using the optical depth τ , with extinction coefficient $\sigma_t = \sigma_s + \sigma_a$, defined as the sum of scattering (σ_s) and absorption (σ_a) coefficients:

$$\tau(x, x') = \int_x^{x'} \sigma_t(t) dt \quad (4)$$

Note that, in contrast to out-scattering, absorption and in-scattering, emission was omitted in the rendering equation for simplicity. Since the rendering equation cannot be computed analytically, solving the integral numerically would involve sampling the function at many distances, each with many directions. Additionally, L_i must be computed with the same rendering equation to allow multiple scatter events. Since this would be computationally too complex, the Monte Carlo method allows us to compute the radiance at random positions along the ray with light being in-scattered from random directions. By averaging many of such Monte Carlo samples into a single image we can progressively generate a smooth final result. By means of multiple sampling, the convergence of the method can be accelerated considerably.

The medical data is illuminated using image-based lighting by high-dynamic range lighting environments, which can either be captured photographically or generated synthetically. Photographically captured lighting leads to a very natural appearance of the data when compared to images created using the traditional ray casting method. Such natural lighting in combination with the accurate simulation of photon scattering and absorption, leads to photorealistic images (see Fig. 1) that resemble many shading effects that can be observed in nature, such as soft shadows, ambient occlusion, volumetric scattering and subsurface photon interaction. By modelling a virtual camera with variable aperture, focal length and exposure, additional effects such as depth-of-field and motion blur can be produced. Motion blur allows movies generated using our key frame animation engine to be smoother during fast camera movements while, similar to photography, depth-of-field effects allow to focus the attention of a viewer on a particular structures.

Beyond photorealism the algorithm also permits to visualize invisible or hidden processes such as the propagation of electrical activation on the heart surface or metabolic processes in the body. Such hyper-realistic images are created by modelling visible light photon emission from voxels affected by electrical activation, increased metabolism indicated by Positron Emission Tomography (PET) or the detection of chemical compounds such as monosodium urate from a dual-energy CT scan (Fig. 2).

The combination of different imaging modalities in a single picture, such as PET, MR and CT as well as simulated and computed



Fig. 2. Gout visualization. Modelling photon emission from urate detection by a dual-energy CT scan.

data provides important flexibility to show the spatial relation of anatomical structures and functional data (see Fig. 3).

Another important application of Cinematic Rendering is the visualization of dynamic processes from 4D CT or MR scans in combination with time-dependent data from simulations. All such data sources can be combined frame-by-frame and played using an animation engine to create photorealistic movies which allow conveying an effective clinical message to the target audience.

While diagnostics will certainly still rely on traditional planar reconstruction based visualization methods, we have strong indications that special diagnostic applications might benefit from the flexibility and expressiveness of the new Cinematic Rendering technology. For instance, a robust demand for such visualization methods can be seen for surgery planning and intraoperative imaging, where a good spatial understanding of the anatomy and processes in the human body is required (Fig. 4).

Furthermore, the use of hyper-realistic imaging for anatomical education of medical students as well as the general public is obvious. And finally, such images are ideal for the efficient communication of findings, diagnoses and surgery results, either among medical professionals or to patients, potentially increasing trust in clinical decisions and procedures.

3. Artificial intelligence and image understanding

Handling the complexity of medical images involves understanding thousands of anatomical classes and concepts, while numerous relationships are necessary to symbolically represent the phenotypic structure of the human body (FMA, 2012). There are multiple dimensions along which this information can be structured, for example one can look at the human anatomy from a regional point of view (limbs, head), or constitutional point of view (lymphatic duct, skin) or system (nervous, cardiovascular, musculoskeletal). Such ontology-based or symbolic representation is a form that is understandable by humans and it is also navigable, parseable and interpretable by machine-based systems.

Fast and robust anatomical concept extraction is a fundamental task in medical image analysis that supports the entire workflow from diagnosis, patient stratification, therapy planning, intervention and follow-up. Current state-of-the-art solutions are based on machine learning, being enabled by the availability of large annotated medical databases and the increased computational capabilities (Zheng and Comaniciu, 2014). Typical methods use example images of the anatomy of interest to learn a classifier that will be able to discriminate between inputs that contain the target anatomy or something else. Such classifiers can be used to automatically label images, detect landmarks or segment the target object (see Fig. 5).

For example, in the context of object detection, the classifier is scanned over all possible values of the parameter space (say translation T , rotation R and scale S) to find the high probability regions that will correspond to object location. This is done by using a classifier that will approximate the probability $p(T, R, S|I)$ for an image I , where the classifier is trained with object image features for one class and non-object image features for the other class. At runtime the object is determined by regions of the parameter space $(\hat{T}, \hat{R}, \hat{S})$ with high probability: $\operatorname{argmax}_{T, R, S} p(T, R, S|I) = (\hat{T}, \hat{R}, \hat{S})$.

In the past years, we have developed technologies such as Marginal Space Learning (MSL) (Zheng and Comaniciu, 2014) for automated and efficient scanning of medical images. The MSL method consists in learning image-based classifiers in high probability marginal spaces of the object parameterization. A set of classifiers are trained in stages: first in the translation hypotheses space $\Omega_T(I)$, where I is the image, next the augmented translation-orientation hypotheses space $\Omega_{T, R}(I)$ and finally in




Fig. 3. Human brain visualization. Left: Cinematic Rendering of a Magnetic Resonance (MR) image of the brain acquired with a 7T scanner. Data courtesy of Max Planck Institute, Leipzig, Germany. Right: Cinematic Rendering of three anatomical MR slices, functional MR Imaging (fMRI), and fiber data computed from Diffusion Tensor Imaging (DTI) of the brain. The activation of the speech center captured by fMRI is modelled using light emission and results in the yellow glowing lighting effect on the anatomy.



Fig. 4. Polytrauma visualization. Cinematic Rendering of a polytrauma patient with multiple spinal and costal fractures. CT image data courtesy of Vancouver General Hospital, Canada.




Fig. 5. Identifying and segmenting anatomical structures in a whole body CT scan (only partial views are shown). Tens of anatomical landmarks are being identified and multiple organs such as heart, lungs, liver, kidneys, spleen, prostate and bladder are segmented and quantified (Seifert 2009).

the translation-orientation-scale hypotheses space $\Omega_{T,R,S}(I)$. The spaces are constructed by augmenting the high probability hypotheses with all the possible discrete values of the next parameter space such as:

$$\begin{aligned} \operatorname{argmax}_T p(\Omega_T|I) &= \Omega_{\hat{T}} \xrightarrow{\text{Augment with discrete rotations}} \Omega_{\hat{T},R} \\ \operatorname{argmax}_{T,R} p(\Omega_{\hat{T},R}|I) &= \Omega_{\hat{T},\hat{R}} \xrightarrow{\text{Augment with discrete scales}} \Omega_{\hat{T},\hat{R},S} \\ \operatorname{argmax}_{T,R,S} p(\Omega_{\hat{T},\hat{R},S}|I) &= \Omega_{\hat{T},\hat{R},\hat{S}} \end{aligned} \quad (5)$$

The MSL technique is generic and can be extended to any type of parameterized spaces. We used MSL to automatically determine hundreds of landmarks, segment, track and quantify all main organs, delineate and index the vascular tree, brain structures and the skeleton.

Recent advances in machine learning and artificial intelligence have created end-to-end learning architectures where all stages of the processing are jointly optimized. For example, representation learning with Deep Neural Networks (DNN) enables automatic extraction of representative image features without the need of feature engineering (Zheng 2015). DNN allow learning complex patterns from very large heterogeneous image databases. We have recently introduced Marginal Space Deep Learning (MSDL) (Ghesu 2016a) that combines the strength of automated feature design of DNN with efficient learning in marginal spaces.

In MSDL (Fig. 6) the classifier is trained directly on parameterized image patches and used to estimate the probability distribution: $\mathcal{R}(\Omega(I); w, b) \approx p(\Omega(I)|I)$ where \mathcal{R} is a deep neural network response function parameterized by the weights w and biases b of each layer. In addition, more efficient scanning of DL networks is achieved through network approximation (sparsification) techniques by minimizing the residual $\|\mathcal{R}(\Omega(I); w, b) - \mathcal{R}(\Omega(I); w_s, b_s)\|$, where w_s and b_s are the weight and bias of the approximated sparse network $\mathcal{R}(\Omega(I); w_s, b_s)$. As a result, $\mathcal{R}(\Omega(I); w_s, b_s)$ has much fewer parameters and/or access much less data from the image making possible scanning for parameterized objects in 3D or 4D images. With MSDL, we have shown significant performance improvements in terms of both accuracy and speed on aortic valve detection in volumetric ultrasound and landmark detection in CT scans (Ghesu 2016a).

With the goal to add more intelligence into image analysis, our recent work has focused on artificial intelligence agents that can be trained using Deep Reinforcement Learning (DRL) techniques to simultaneously model both the object appearance and the object search strategy as a unified behavior (Ghesu 2016b). The idea is to train an agent that can navigate within an image to find an anatomy of interest. In other words, the agent learns automatically optimal paths that converge to the target object, thus eliminating the need for exhaustive search (Ghesu 2016b).

A typical Reinforcement Learning (RL) technique is modeled as a Markov Decision Process (MDP) defined on the tuple $(S, \mathcal{A}, Tr, r, \gamma)$, where S represents a finite set of agent states, \mathcal{A} represents a finite set of actions that the agent can perform to interact with the environment, $Tr: S \times \mathcal{A} \times S \rightarrow [0,1]$ is a stochastic transition function between two states by performing a specific action, $r: S \times \mathcal{A} \times S \rightarrow \mathbb{R}$ is the scalar reward expected after a state transition and γ is a future rewards discount factor. One target in RL is to find an optimal of the action-value function $Q^*(s, a): S \times \mathcal{A} \rightarrow \mathbb{R}$ that corresponds to the maximum expected future rewards when performing action a in state s : $Q^*(s, a) = \max_{\pi} E[r_t | s_t = s, a_t = a, \pi]$ where π represent an action policy that determines the behavior of the agent. For object detection, such an agent can be trained by having S as the current estimates of the object parameters given the image (e.g. spatial coordinates), \mathcal{A} the discrete steps of parameter changes and a reward system that is related to how close the agent gets to the target by performing the actions. Given the model definition, a DNN can be trained to approximate the optimal action-value function Q_* directly from the image values parameterized by the current state or object parameters. The optimal action-value function implicitly defines the optimal policy π_* which guides the agent in finding the target object. This paradigm, where the agent simultaneously learns an object model and how to use the

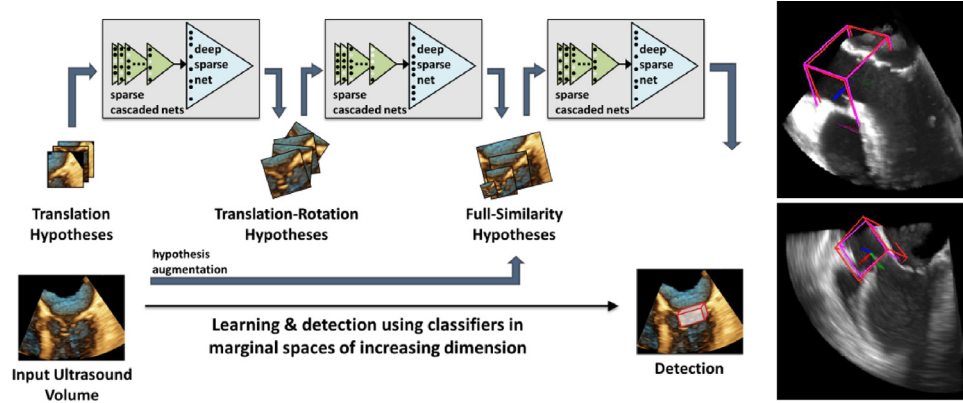


Fig. 6. Marginal Space Deep Learning. Left: Learning in increasingly dimensional spaces focused on high probability regions with deep neural sparse networks. Right: Example of aortic-valve detection in volumetric ultrasound.

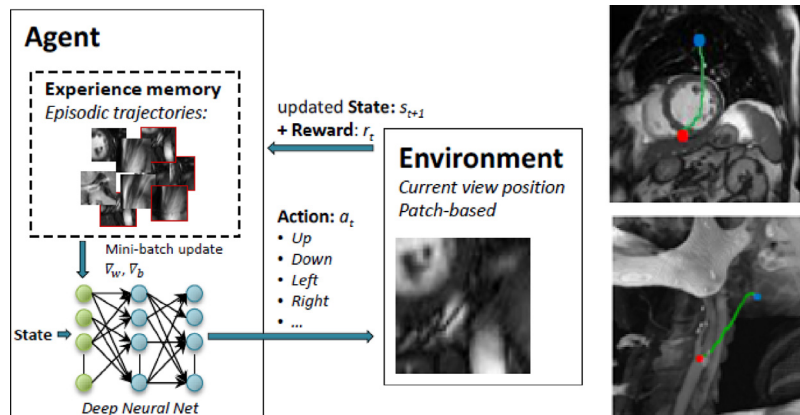


Fig. 7. Artificial agent for landmark detection. Left: Interaction of an artificial agent with the environment for detecting anatomical landmarks: Current state is defined by the image window and the agent performs the optimal action according to the learned behavior, which results in a new state and reward feedback. Right: 2D and 3D paths for detecting anatomical landmarks (in blue: starting point of the agent, red: target).

model, can be extended to a wide variety of image parsing actions (Fig. 7).

Combining these types of learning techniques and classifiers with ontology based representation allows for semantic navigation of all the available patient data (Seifert 2011). In particular, expanding end-to-end learning systems to include medical knowledge with powerful learning-based representation and reasoning systems will allow building hierarchical representations dynamically based on the current task. This will facilitate comprehensive automated analysis and reporting based on integrated imaging and non-imaging information, past reports and embedded medical knowledge. For instance, these methods will also enable integrated analysis of patient history with support for semantic search and case comparison by finding similar cases and treatments with related clinical knowledge and guidelines. Finally, the method facilitates knowledge sharing and population analytics. Such information empowers the radiologist towards increased efficiency and reduced uncertainty.

4. Support for minimally-invasive procedures

Progress in medical imaging technologies and image analysis are making increasingly complex minimally invasive procedures possible. Techniques like heart valve repair or replacement can now be performed percutaneously, a relevant example being Transcatheter Aortic Valve Replacement (TAVR). Effective execution of minimally invasive procedures strongly relies on medical imaging.

First, the devices need to be selected to fit patient's anatomy. To that end, quantitative imaging is used to accurately measure the anatomy under consideration. For instance, the size of TAVR devices is determined from the dimensions of the aortic root. This is usually performed on CT data, but for patients suffering from kidney failure, novel, full-volume, real-time 3D TEE (trans-esophageal echocardiography) now enables the quantification of heart valves and blood flow without contrast agent. Fully exploiting this new imaging modality, we recently developed an advanced machine learning technology to estimate a personalized model of the heart valves (Fig. 8). In brief, the algorithm, based on Marginal Space Learning, first detects the region of interest (ROI) where the valve is located. Within that ROI, key anatomical landmarks are detected and a parameterized triangulated surface is fitted to model the patient's valve. The detectors of each stage of the algorithm are trained from a large database of annotated images. The method is generic, and has been applied to other imaging modalities, like CT (Ionasec 2010).

Second, the operator needs image guidance to effectively deliver the device. Angiography is the modality of choice in the hybrid OR as it allows the visualization of catheters in real-time. However, soft tissues are hard to distinguish in these images. Advanced navigation concepts based on augmented reality have thus been investigated to enhance the angiography images with overlay of the targeted organ. 3D preoperative images are registered to the 2D scene (often facilitated by the injection of contrast) using pattern matching, multi-organ registration or multi-view reconstruction. Supported by imaging technologies, new minimally invasive




Fig. 8. Real-time 3D TEE and 3D Doppler. New TEE imaging allows full-volume, real-time visualization of cardiac anatomy (B-mode) and blood flow (color Doppler). For the first time, the clinician can visualize jointly the valve anatomy and potential insufficiency.




Fig. 9. 3D TEE – Angiography fusion. Left: Real-time overlay on angiography image of the valve model estimated from 3D TEE image acquired at the same time (see Fig. 8); Right: Real-time overlay of the 3D Doppler generated by the 3D TEE probe.

procedures are emerging, like the recent transcatheter mitral valve replacement repairs techniques, which in turns require unprecedented levels of registration accuracy and robustness to cope with moving devices and image artifacts.

The future will likely go towards real-time 3D guidance, as hybrid solutions combining angiography and 3D TEE are becoming available. Through TEE probe pose estimation in the 3D space from angiography images, real-time ultrasound images can be automatically registered to the angiography space. User-defined landmarks, anatomical models or the ultrasound images directly can then be overlaid to the angiography image to guide the cardiologist towards the target (Fig. 9). To reach the accuracy and speed requirements for real-time intervention guidance, we recently introduced a 3D TEE probe pose estimation based on deep learning (Miao 2016), yielding high accuracy at a frame-rate of 15fps.

5. Decision support through patient specific computational models

While increasingly robust and accurate quantification methods are being available, new solutions based on computational models of human physiology are being investigated to extract more physiological information from the images and facilitate patient-specific planning through predictive algorithms.

A first example is Fractional Flow Reserve (FFR), the current gold standard parameter that characterizes coronary stenosis severity. In standard of care, FFR is measured invasively using pressure catheters. During the past years, we have developed non-invasive, image-based FFR methods ($cFFR_{CFD}$) based on CT images and reduced-order computational fluid dynamics (CFD) models, making it possible to calculate FFR at the bed-side (Sharma 2012) with excellent performance. Furthermore, with the advances in machine learning, we have demonstrated that it is now possible to calculate FFR non-invasively in seconds on a standard workstation. Based on deep learning, the new approach consists in learning the CFD model directly from anatomical features (Itu 2016). First, a database of 12,000 coronary geometries with more than 1000,000 coronary segments was computed, with randomly positioned stenosis. Second, a reduced-order CFD model was used on all 12,000 geometries to calculate the resulting FFR. Both heart and systemic circulation models were included for proper boundary conditions. Third, a deep-network ($cFFR_{ML}$) was trained to predict the FFR value given geometric features computed upstream, at and downstream the stenosis. Tested on 127 unseen lesions from 87 patients, $cFFR_{ML}$ could be calculated in 2.4 s in average, with a correlation coefficient of 0.9994 ($p < 0.001$) and no bias with respect to $cFFR_{CFD}$. Compared to the invasive FFR value, $cFFR_{ML}$ sensitivity was 81.6%, specificity 83.9% and accuracy 83.2%, achieving similar performance as other non-invasive FFR methods (Fig. 10).




Fig. 10. Testing deep learning based FFR. Left: $cFFR_{CFD}$ computed using reduced order CFD model; Right: $cFFR_{ML}$ values. The deep learning approach could capture the FFR values accurately, while reducing the calculations to seconds.

Similarly, we have developed machine learning techniques to estimate a patient-specific model of cardiac electrophysiology (EP) for cardiac resynchronization therapy (CRT) planning (Zettinig 2014; Kayvanpour 2015). In particular, in (Neumann 2016), we introduce an intelligent agent trained following RL concepts to estimate cardiac electrical conductivities directly from 12-lead ECG. Following the notations introduced in Section 2, we defined the RL Markov Decision Process $(S, \mathcal{A}, Tr, r, \gamma)$ as follows. The states $s \in S$ are discretized objective values (absolute difference between target and current QRS duration and electrical axis). The actions $a \in \mathcal{A}$ are to increment and decrement the left endocardial, right endocardial, and myocardial conduction velocity (six actions in total). In a first stage, the agent learns through random exploratory simulations the state transition probabilities Tr , which encode how the model behaves when the electrical conductivities change. Given this knowledge, the best personalization strategy is learned through RL. The reward r is defined such that at every step t , starting from a state s_t , the agent receives a negative reward for all actions $a_t \in \mathcal{A}$ except for the one a_{t*} that leads it to the parameters for which the EP model best matches the observed ECG. Finally, the parameter γ was set to 0.9 or higher to favor long-term rewards (and thus global optimum). Tested on 83 consecutive patients, the artificial agent could achieve similar goodness of fit as a hand-crafted, state-of-the-art optimization method (Neumann 2016), while being 2.5 times faster.

Such personalized models can be used for patient-specific therapy planning. For instance, a user could use the model to virtually test different CRT pacing protocols (Kayvanpour 2015). She would virtually place the CRT leads, program the virtual device, and update the model to visualize the impact of the CRT pacing on cardiac function (Fig. 11), thus increasing the confidence on the therapy and the best strategy to apply to the patient under consideration.

6. Conclusion

In this paper we discussed recent technologies that will most likely make an important impact on medical imaging. Techniques like Cinematic Rendering will help increasing the sensitivity and specificity of images, by enhancing the pathology conspicuity. Advanced image understanding will streamline the image measurements and image interpretation, by increasing the speed of reading, while introducing more reproducibility in the system. The new




Fig. 11. Personalized CRT planning. Effect of ventricular pacing on cardiac electrophysiology as calculated by an individualized computational model of heart function.

heart valve technologies, based for instance on 3D TEE/Doppler, will help a more precise characterization of the patient's anatomy in the OR, while the 3D TEE – Angiography Fusion will support better guidance. Finally, patient specific computational modelling opens the door to a new generation of decision support systems that help clinical decision making not only by integrating and analyzing data from different sources, but also by modelling both the anatomy and function of the patient, thus exhibiting enhanced predictive power.

Disclaimer

This feature is based on research and is not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.

Acknowledgement

The authors would like to thank all colleagues at the Siemens Healthcare Technology Center who contributed to the works discussed in this article. We would also like to thank our clinical and non-clinical collaborators for the strong support over many years.

References

- Engel, K., 2016. Real-Time Monte-Carlo Path Tracing of Medical Volume Data, GPU Technology Conference, APRIL 4–7, 2016. San Jose Convention Center.
- Ghesu, F.C., Krubasik, E., Georgescu, B., Singh, V., Zheng, Y., Hornegger, J., Comaniciu, D., 2016a. Marginal space deep learning: efficient architecture for volumetric image parsing. *IEEE Trans. Med. Imaging*.
- Ghesu, F.C., Georgescu, B., Mansi, T., Neumann, D., Hornegger, J., Comaniciu, D., 2016b. An artificial agent for anatomical landmark detection in medical images. *Medical Image Computing and Computer Assisted Intervention (MICCAI)*.
- FMA, Foundational Model of Anatomy Ontology. 2012. https://en.wikipedia.org/wiki/Foundational_Model_of_Anatomy
- Ionasec, R.I., Voigt, I., Georgescu, B., Wang, Y., Houle, H., Vega-Higuera, F., Navab, N., Comaniciu, D., 2010. Patient-specific modeling and quantification of the aortic and mitral valves from 4-d cardiac CT and TEE. *IEEE Trans. Med. Imaging* 29, 1636–1651.
- Itu, L., Rapaka, S., Passerini, T., Georgescu, B., Schwemmer, C., Schoebinger, M., Flohr, T., Sharma, P., Comaniciu, D., 2016. A machine learning approach for computation of fractional flow reserve from coronary computed tomography. *J. Appl. Physiol.*
- Kayvanpour, E., Mansi, T., Sedaghat-Hamedani, F., Amr, A., Neumann, D., Georgescu, et al., 2015. Towards personalized cardiology: multi-scale modeling of the failing heart. *PLOS ONE* 10 (7).
- Miao, S., Wang, Z.J., Liao, R., 2016. A CNN Regression Approach for Real-Time 2D/3D Registration. *IEEE Trans. Med. Imaging*.
- Neumann, D., Mansi, T., Itu, L., Georgescu, B., Kayvanpour, E., Sedaghat-Hamedani, AmrA., et al., 2016. A self-taught artificial agent for multi-physics computational model personalization. *Med. Image Anal.*
- Paladini, G., Petkov, K., Paulus, J., Engel, K., 2015. Optimization techniques for cloud based interactive volumetric monte carlo path tracing. *Industrial Talk, EG/VGTC EuroVis*
- Seifert, S., Barbu, A., Zhou, S.K., Liu, D., Feulner, J., Huber, M., Suehling, M., Cavallo, A., Comaniciu, D., 2009. Hierarchical parsing and semantic navigation of full body CT data. *SPIE Med. Imaging*.
- Seifert, S., Thoma, M., Stegmaier, F., Hammon, M., Kramer, M., Huber, M., Kriegel, H.-P., Cavallaro, A., Comaniciu, D., 2011. Combined semantic and similarity search in medical image databases. *SPIE Med. Imaging*.
- Sharma, P., Itu, L., Zheng, X., Kamen, A., Bernhardt, D., Suci, C., Comaniciu, D., 2012. A framework for personalization of coronary flow computations during rest and hyperemia. In: *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pp. 6665–6668.
- Zetting, O., Mansi, T., Neumann, D., Georgescu, B., Rapaka, S., et al., 2014. Data-driven estimation of cardiac electrical diffusivity from 12-lead ECG signals. *Medical Image Anal.* 18 (8), 1361–1376.
- Zheng, Y., Comaniciu, D., 2014. *Marginal Space Learning for Medical Image Analysis: Efficient Detection and Segmentation of Anatomical Structures*. Springer.
- Zheng, Y., Liu, D., Georgescu, B., Nguyen, H., Comaniciu, D., 2015. 3D deep learning for efficient and robust landmark detection in volumetric data. *Medical Image Computing and Computer Assisted Intervention (MICCAI)*.