

Creating a Digital Twin of Spinal Surgery: A Proof of Concept

Jonas Hein^{1,2,*} Frédéric Giraud¹ Lilian Calvet¹ Alexander Schwarz²
 Nicola Alessandro Cavalcanti¹ Sergey Prokudin² Mazda Farshad¹ Siyu Tang²
 Marc Pollefeys² Fabio Carrillo¹ Philipp Fürnstahl¹

¹Balgrist University Hospital, University of Zurich, Zurich, Switzerland

²ETH Zurich, Zurich, Switzerland



Figure 1. Digital photograph of a spinal surgery (left) and rendering of its digital twin (right) obtained using our proof of concept for surgery digitalization.

Abstract

Surgery digitalization is the process of creating a virtual replica of real-world surgery, also referred to as a surgical digital twin (SDT). It has significant applications in various fields such as education and training, surgical planning, and automation of surgical tasks. In addition, SDTs are an ideal foundation for machine learning methods, enabling the automatic generation of training data. In this paper, we present a proof of concept (PoC) for surgery digitalization that is applied to an ex-vivo spinal surgery. The proposed digitalization focuses on the acquisition and modelling of the geometry and appearance of the entire surgical scene. We employ five RGB-D cameras for dynamic 3D reconstruction of the surgeon, a high-end camera for 3D reconstruction of the anatomy, an infrared stereo camera for surgical instrument tracking, and a laser scanner for 3D reconstruction of the operating room and data fusion. We justify the proposed methodology, discuss the challenges faced and further extensions of our prototype. While our PoC partially relies on manual data curation, its high quality and great potential motivate the development of automated methods for the creation of SDTs.

1. Introduction

Surgery digitalization is the process of creating a virtual replica of a real-world surgery, most commonly known as a surgical digital twin (SDT). The digital twin concept was first introduced by [14] and consists of three main components: a physical object or process along with its environment, its digital replica, and the data and communication links that connect the physical and digital entities. It is a specific application of digitalization, namely the process of converting information from a physical format to a digital one, focusing on the replication and simulation of physical entities. One of the main objectives of a SDT is the high-fidelity representation of relevant entities and their interactions during the surgery, including the patient, surgical instruments and devices, as well as medical staff.

Surgery digitalization has diverse downstream applications [37], ranging from optimizing education and enhancing the capabilities of surgical services to enabling the training of surgical robots [48]. In the realm of education, surgery digitalization has the potential to provide medical

*jonas.hein@inf.ethz.ch

students and surgeons with realistic and interactive virtual environments to practice surgical techniques and understand human anatomy without the need for real anatomical models, which are often expensive and scarce, limiting the hands-on learning experience [34]. It may facilitate operative performance assessment, formative feedback and surgical credentialing by avoiding the need for manual review and assessment of surgical videos [35, 42]. The ability of replaying or streaming a surgery may enable novel use-cases in the areas of quality control and remote surgery [23]. In the context of workflow optimization, surgery digitization enables the optimization of resource allocations through machine learning (ML)-based surgical phase recognition [18], or the automatic generation of surgery reports [27]. In surgical navigation and robotic surgery, SDTs can help to reduce the sim-to-real gap by providing accurate and realistic environments in which robots and ML-based applications can be trained before being deployed in the real world [2, 9, 15, 19].

Surgery digitalization necessitates the fusion of available information from sensors and prior knowledge into a common spatio-temporal representation that accurately describes the state of its physical twin [12]. Data may come from different modalities, including imaging, sound and text. To date, imaging technologies are yet predominant. For the acquisition of information associated with anatomy, medical imaging technologies such as computed tomography (CT), magnetic resonance imaging (MRI), ultrasound, fluoroscopy and endoscopy are preferred. Optical cameras remain the solution of choice for other components of the surgery, namely for capturing information associated to medical staff, the operating room (OR) and its devices, and surgical instruments [20, 21]. They are widely used due to their ability to capture detailed visual information in a non-invasive manner while being able to record events in real-time. Once collected, the data is processed to create a model whose complexity varies depending on the type of information to be encoded and on the downstream applications. A low level representation may consist of a raw multi-view RGB(-D) video [21] while a high level representation may consist of a detailed semantic and geometric representation of the surgical scene [18, 41]. High level representations could also include advanced behavioral models such as those proposed by [50].

In this work, we describe a proof of concept (PoC) to digitize a segment of spine surgery, more specifically the pedicle screw drilling done within the pedicle screw placement procedure, in near-realistic surgery conditions. It deviates from a real surgery in that it is performed by a single surgeon without assistance to mitigate occlusions and limit the number of cameras needed. The surgery is performed ex-vivo on a human specimen in an operating room dedi-

cated to translational research in surgery², enabling an extensive data collection that would be infeasible during real patient treatment. The specific problem being addressed is how to combine cutting-edge 3D scanning technologies with optimal data fusion and modelling techniques to create a spatio-temporal 3D model of a surgical scene that verifies the following four criteria: it must be (C1) *faithful with respect to geometry*, meaning that the dimensions and spatial relationships in the model should accurately reflect those of the actual surgical setting over time, (C2) *explicit*, (C3) *modular*, which means it is built from smaller, distinct components that represent real-world objects within the surgery, and (C4) *complete*, encompassing the entire surgical scene to provide a full and uninterrupted representation. Criterion C1 ensures the model supports highly immersive training and education for surgery. Additionally, it enables precise 3D measurements, essential for surgical navigation, planning and quality control. Criterion C2 guarantees the model is interpretable, allows for measurements using standard metrics, and is compatible with widely used rendering engines. Parametric representations, especially for the medical staff and instrument's locations should be prioritized. Criterion C3 allows for the individual manipulation of different components (anatomy, surgeon, surgical instruments, etc.), enabling customizable simulations and object level reasoning in the context of surgical workflow or activity recognition. Finally, criterion C4 guarantees a holistic representation of the surgery, from which every downstream application can benefit. However, this requirement necessitates dedicated data acquisition setups, as the data collected during surgeries today is still too sparse to provide such a holistic representation.

In response to these criteria, we contribute a surgery digitalization approach which generates a SDT as a set of textured 3D meshes, representing the furnished OR and the anatomy (static rigid), the surgeon and the surgical drill (dynamic), in a shared spatio-temporal representation. The choice of 3D scanning technologies being used and the data fusion and modeling processes are described. The obtained SDT is made publicly available. Although our PoC partially relies on manual data curation and assumptions that still diverge from actual surgical settings, it is expected to motivate the development of fully automated and functional methods for surgery digitalization under real surgical conditions.

The remainder of this work is organized as follows. The state of the art is discussed in Section 2. The proposed methodology for surgery digitalization is detailed in Section 3. Section 4 presents a quantitative and qualitative evaluation of the quality of our SDT. Section 5 discusses the proposed methodology and its limitations, perspectives and potential applications, before our conclusions are drawn in Section 6.

²<https://www.or-x.ch/>

2. Related work

Surgical data science Over the last two decades the field of surgical data science emerged from the need for systematically captured and structured medical data to improve the quality of interventional healthcare [32]. The importance of this field further increased with the rapid advance of ML methods in the last decade. State-of-the-art deep learning methods typically require large amounts of structured training data, and the lack thereof is one of the major obstacles in the field [33]. However, patient-related data is still not systematically recorded and stored. In high-income countries, which benefit from access to advanced healthcare systems and robust IT infrastructures, difficulties arise from navigating regulatory and policy frameworks, as well as from the high complexity of medical data [16, 33]. Relevant information is often distributed across several disconnected systems and in different data modalities which also makes data collection in a standardized and systematic way highly challenging. The digitalization of surgeries would enable a more standardized and structured data collection process.

Surgical digital twin Digital twins aim to be a perfect virtual representation of their physical counterpart, such that observations of the digital twin yield the same information as observations of the physical one. Later works extended this original definition by [14] to include physical, bio-mechanical, or behavioral models that enable the simulation, prediction of future states, and closed-loop optimization of task-specific objectives [4]. In the medical field, a digital twin of the patient has the potential to enable patient-specific optimal treatment [29]. Previous works have proposed digital twins for specific anatomies and interventions, including knee arthroscopy and skull base surgery [4, 44]. Most approaches rely on the registration of a preoperative 3D model of the patient-specific anatomy with the patient and the surgical instruments, typically through marker-based tracking. However - to the best of our knowledge - no existing model aims to capture a full surgery yet. On the scale of an operating room, the interactions between patient, instruments, surgeons and medical personnel are highly relevant for an accurate description of the current state of the surgery. Surgical scene graphs [41] are a lightweight representation of high-level spatial and semantic relationships of entities in the OR. Several works proposed to estimate surgical scene graphs from video [52]. Similarly, surgical process models [40] have been proposed to hierarchically describe the surgical phases and steps comprising an intervention. While these graph-based representations may be beneficial for high-level tasks such as visual question answering [49] or surgical phase recognition [18], they abstract the low-level geometry required for a high-fidelity representation of the surgical scene.

3D reconstruction Various types of technologies have been developed to digitize the 3D structure of the physical world


with high fidelity. Laser scanning and structured light scanning have established themselves as powerful tools for their direct and precise acquisition of 3D data. Laser scanning, utilizing the principle of light detection and ranging (LIDAR), offers unparalleled accuracy in capturing large-scale environments and intricate details over vast distances. On the other hand, structured light scanning, which projects patterned light onto objects and measures deformations through a camera system, excels in capturing high-resolution surface details of smaller objects within controlled environments. In clinical research, both LIDAR and structured-light approaches have been explored for the 3D reconstruction of soft tissue [3, 5, 11, 36]. While these methods provide robust solutions for 3D data acquisition, they present limitations in terms of equipment cost, operational complexity, and environmental constraints, which can hinder their applicability in diverse scenarios.

In contrast, photogrammetry, which reconstructs 3D models from 2D optical images, is an alternative that offers versatility and accessibility beyond the capabilities of laser or structured-light scanning. The literature on computer vision presents photogrammetry methods that are capable of deducing both the shape and appearance of objects from a collection of uncalibrated optical images. Recent Neural Radiance Fields (NeRFs) [39], their extensions to surface reconstruction [46] and large scale acquisitions [25], and even more recently Gaussian splatting techniques [22], have demonstrated remarkable performance in both controlled and uncontrolled environments. These advances enable the creation of highly detailed and photorealistic renderings from relatively sparse image datasets, marking a significant leap forward in the field's capabilities. Recent works applying these techniques on medical imaging data have obtained impressive results [13, 28], however there remain several challenges in their practical application, such as the handling of dynamics and long temporal sequences, or the compatibility with standard rendering engines.



Pose estimation Most of the above-mentioned computer vision methods assume a mainly rigid scene. However, a SDT setting includes dynamic objects, primarily medical staff, surgical instruments, and the anatomy. To estimate the pose of surgical instruments, marker-based navigation systems like the *FusionTrack* (Atracsys LLC, Puidoux, Switzerland), which combine a stereo-camera with infrared-sensitive markers mounted on instruments, show sub-millimeter accuracy and remain the gold standard solution. Their main limitations are the line of sight issue and limited working volume, which have motivated the development of marker-less tracking approaches [10, 17]. To estimate the pose of the medical staff, motion capture systems can be used. Vicon systems³ are the gold standard technology for motion capture in film and video game pro-

³<https://www.vicon.com/>

Single laser scan:



Fused laser scans:



● Markers ○ Scanning position




Figure 2. Generation of the reference point cloud from multiple laser scans. The first row shows the point cloud obtained from a single laser scan, illustrating the occlusion challenge. In comparison, the bottom row shows the reference point cloud after fusing all 8 scans. The top view in the center indicates the 21 marker locations and 8 scanning positions within the room. We also indicate the origin of the reference frame, which lies in the ground plane.

duction. They are also widely used in sports biomechanics and virtual reality applications. Multiple high-speed cameras placed around a controlled environment are employed to track reflective markers attached to the subject. These systems are costly and the required amount of cameras remains invasive in the context of surgery digitalization. Other motion capture systems such as the XSens⁴ rely on inertial measurement units (IMUs). However, these systems are impractical for routine captures, largely due to their time-consuming setup and calibration processes. An easy-to-use alternative are marker-less body pose estimation methods, which have been developed based on computer vision techniques [6]. Similar to marker-based systems, they produce a skeletal representation of the body from image data. This skeletal representation seamlessly integrates with parametric models for the human body [30] and hands [43], whose parameters optimally explain the image content while facilitating the creation of a surface representation of the body as a dynamic mesh. Such marker-less pose estimation approaches enable non-invasive data acquisition setups, which is highly relevant for dynamic and restrictive environments such as ORs.

3. Methodology

In this section, we describe our prototype for surgery digitalization. Our data acquisition setup comprises five RGB-D cameras for dynamic 3D reconstruction of the surgeon, a high-end camera for 3D reconstruction of the anatomy, and an infrared (IR) stereo camera for surgical instrument tracking. We additionally employ a laser scanner for 3D reconstruction of the OR and its devices, and for the

⁴<https://www.movella.com/products/xsens/>

fusion of all captured entities in a shared reference frame. We first describe the acquisition and fusion of data associated with static elements and their modelling in Sections 3.1 and 3.2, and those associated with dynamic elements in Section 3.3.

3.1. Reference frame acquisition

The basis of our SDT is a 3D representation of the OR with metric scale, which serves as our reference frame for the registration of all static and dynamic elements. We employed a *Faro Focus 3D 120* laser scanner (FARO Technologies Inc., Lake Mary, FL, USA) to generate a point cloud representation of the room. To minimize occlusions, we conducted 8 scans from various positions, which were subsequently fused. To this end, we temporarily and uniformly positioned 21 markers throughout the space, as depicted in Figure 2. These markers were used as point primitives for a point-to-point registration of all 8 point clouds. Finally, the origin of the reference point cloud was established at the center of the floor. Its orientation was defined by the first two main components from principal component analysis applied to the floor points.

In this PoC we assume that the ceiling objects, OR equipment, and the instrumentation table are static. Based on these assumptions, this reference point cloud is utilized to integrate all components of the model.

3.2. Modeling the operating room

The objective of this phase was to create a detailed and visually accurate virtual model of the operating theatre, including permanently mounted devices like the OR lamps and displays. For this purpose, we utilized the open-source 3D modeling software *Blender* (Stichting Blender Foun-

dation, Amsterdam, Netherlands) in conjunction with its Eevee rendering engine. The CAD models of the room and ceiling elements were crafted by a professional graphics artist, while utilizing the reference point cloud to accurately determine the dimensions. We modeled the textures and materials based on detailed photographs to enhance the visual realism. Given the necessity for this simulation to accommodate various configurations of the OR, the ceiling elements in the model were designed with movable joints, which replicate the kinematics of their real-world counterparts. Additionally, we incorporated the functionality to adjust the parameters of both the ceiling and OR lighting, as well as the contents of the display screens within the OR.

The CAD model of the operating theatre is precisely aligned with the reference point cloud. In this alignment process, the joints of the ceiling objects are manually adjusted to match the reference point cloud. Exemplary renderings can be seen in Figure 5. Although the manual modeling of the operating room is time-consuming, this one-time effort yields a detailed model of an OR which can be the basis for realistic training simulators and synthetic data generation.

Operating table and anatomy We employed a photogrammetry approach to reconstruct the operating and instrumentation tables, as well as the visible surface of the anatomy. For this, we utilized a *Sony Alpha7R* digital single-lens reflex camera (Sony Group Corporation, Tokyo, Japan) to capture 102 sets of images from different viewpoints. These photos were captured just before the start of the surgery and with the anatomy and instruments already placed on the tables. Each set included a focus bracket of five pictures to capture finer details of the tools and spine. Focus stacking was performed using the publicly available code [1] applied to the captured photographs. The commercial photogrammetry software *RealityCapture*⁵ was then used to produce a textured 3D model of the scene from the focus-stacked images. The use of this software was mainly motivated by the quality of its 3D reconstructions that is competitive with state of the art in computer vision, its good compromise between reconstruction quality and computation time, and its camera self-calibration. The focus-stacked images lead to a very fine detailed texture the 3D model benefits from. Reconstruction artifacts were manually removed. The feet of both tables were modeled by hand, as their reconstruction was incomplete due to very challenging surface material and occlusions. The obtained 3D models were then manually aligned with the reference point cloud to accurately reflect the real-world setup.

We furthermore integrated a 3D model of the inner anatomy into our SDT as to capture the interactions between the surgical instruments and the anatomy. To this end, we manually registered the 3D spine model to the vis-

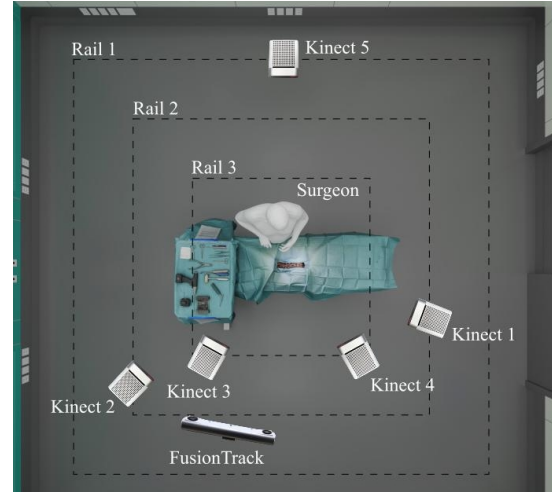


Figure 3. Schematic overview of the experimental setup. Five ceiling-mounted *Azure Kinect* RGB-D cameras capture the motion of the surgeon. A *FusionTrack 500* marker-based tracking system captured the trajectories of the surgical instruments.

ible anatomy surface included in the photogrammetric reconstruction. Due to time constraints we utilized a generic 3D spine model, however this generic model can be easily replaced with a patient-specific preoperative model. These models can be obtained from CT or MRI and are readily available for most orthopedic interventions.

3.3. Motion capture setup

To capture the dynamics of the scene, we deploy a motion capture setup comprising five ceiling-mounted *Azure Kinect* RGB-D cameras (Microsoft Corporation, Redmond, WA, USA) and a *FusionTrack 500* marker-based tracking system (Atracsys LLC, Puidoux, Switzerland). We place four cameras opposite of the surgeon to capture their interaction with the instruments and patient, as shown in Figure 3. These cameras are mounted at different distances to the operating table, such that both the surgical near field and far field are captured. As the surgeon’s lower body is occluded by the operating table in all four cameras, we mount a fifth camera behind the surgeon to complement the four frontal viewpoints and simplify the body pose estimation task. All cameras are mounted above the surgeon’s head height to minimise the intrusiveness of our setup. Due to weight limitations of the camera arms, the tracking system is placed on a tripod. All RGB-D cameras are hardware-synchronized. We follow the approach proposed by [17] to calibrate the cameras extrinsic parameters and to temporally synchronize them with the tracking system. We then registered the multi-camera setup to the surgical environment by solving the perspective-n-point (PnP) problem between 3D points in the reference point cloud and corresponding 2D pixels in one of the cameras. Hereby, we utilized the same 21 mark-

⁵<https://www.capturingreality.com/>



Figure 4. Comparison of the rendered digital twin with the real camera images. The camera perspectives shown from left to right correspond to the *Kinect* cameras 1-5 as shown in Figure 3. The digital twin was rendered in *Blender* using the *Cycles* engine.

ers that were used to register the laser scans, as described in Section 3.1.

Surgical instruments Spinal instrumentation consists of pre-drilling screw trajectories for pedical screw placement. Hereby, an AR-600 battery-powered drill (Arthrex Inc., Naples, FL, USA) is used along with a drill sleeve (Depuy Synthes, Raynham, MA, USA). The drill is tracked via a marker array comprising five IR reflective hemispheres with a diameter of 3 mm attached to the drill body. Following [17], we obtain a 3D model of the surgical drill and marker array using a high-fidelity 3D scanner (Artec3D, Senningerberg, Luxembourg). We registered the marker array to the 3D model by aligning virtual spheres to each hemisphere using the iterative closest point (ICP) algorithm. We obtain a second 3D scan without any attached markers in order to hide the attached marker in the renderings. Both 3D scans are registered using ICP.

Human pose estimation To recover the surgeon’s body pose, we fit the SMPL-H model [30, 43] to the multi-view RGB images. We detect 2D keypoints for all RGB images via *OpenPose* [6]. These keypoints describe 25 distinct anatomical locations of the human body as well as 21 locations on each hand. To take into account that multiple individuals would be present in a real surgery, we follow a simple heuristic to select the keypoints corresponding to the surgeon by computing the mean for each identified person and selecting the person closest to the operating table. Given the 2D keypoints detected in all images and the camera calibration, we compute the 3D keypoints of the surgeon via triangulation. Lastly, we run a multi-stage optimization algorithm over the 2D and 3D keypoints, 2D bounding boxes, and SMPL-H model parameters and weights to compute the SMPL-H model. The optimization algorithm strongly resembles the one from *EasyMocap*⁶, but we added a moving average smoothness term to reduce jittering.

⁶<https://github.com/zju3dv/EasyMocap>

4. Results

While acknowledging the preliminary nature of this research and the requirement for manual intervention, we provide both quantitative and qualitative results to demonstrate the feasibility, accuracy, and potential benefits of our PoC.

Quantitative results The registration error for the laser scan fusion is reported using the root mean square errors (RMSE) from point-to-point registrations performed between two laser scans, namely source and target ones. One laser scan was used as target for all the registrations and all the other ones as sources. We additionally report the 3D reconstruction accuracy for the OR in terms of chamber distance (CD) over overlapping areas between two laser scans. The per-registration RSME and CD reported in Table 1 verify that the fused reference point cloud is millimeter-accurate on the scale of the operating room. We furthermore evaluate the registration error between the laser scans and the photogrammetry model in terms of one-sided CD from the fused laser scans to the photogrammetry model. The obtained CD is 6.72 mm. These results show that our SDT is generally millimeter-accurate, an accuracy which may be partially attributed to the rigidity assumption that holds for our experimental setup.

Following [17], we evaluate the calibration and temporal synchronization errors of the RGB-D cameras in terms of reprojection errors. The errors are averaged over all frames in the calibration sequence and reported in Table 2. We additionally evaluate the registration of the camera array to the reference point cloud by computing the reprojection errors of the reference markers (as discussed in Section 3.1) into each camera. We obtain an average reprojection error of 1.39 pixels.

Qualitative results To showcase our digital twin, we render a video of the spatio-temporal scene with different visual overlays, which is available in the supplementary material. The video is rendered in *Blender 3.3.1* using the *Eevee* ren-




Figure 5. Exemplary renderings of the operating room including the reconstructed operating table and the surgeon’s estimated body pose.

Table 1. Point-to-point registration errors of the laser scans. We choose the laser scan with most visible markers as the reference and register the remaining 7 scans based on all markers visible whose number is indicated in the first row. We report the RMSE of the registered 3D marker positions as well as chamfer distance (CD) between both point clouds, with an outlier filtering of 0.1 m.

Laser Scan	1	2	3	4	5	6	7	Mean
# Markers	12	13	13	14	12	13	12	12.7
RMSE (mm)	7.81	6.42	6.72	5.79	6.95	8.16	6.03	6.84
CD (mm)	4.47	5.02	4.90	4.08	2.90	3.50	3.81	4.10

Table 2. Reprojection errors after the extrinsics calibration and synchronization of each RGB-D camera to the tracking system. We report the mean and standard deviation of reprojection errors over all frames in the calibration sequence. The camera locations are visualized in Figure 3.

Camera	1	2	3	4	5	Mean
Mean error (px)	0.75	0.40	1.06	1.63	0.39	1.19
Std of errors (px)	0.36	0.29	0.89	1.12	0.38	0.92

dering engine. Subtitles, transitions and side-by-side comparisons are added in post-processing using *DaVinci Resolve* (Blackmagic Design Pty. Ltd, Port Melbourne, Australia).

The point clouds computed for each RGB-D camera are combined and outlier filtering is applied. The filtered point cloud is cropped using manually defined bounding boxes to remove static elements like walls and floor. The cropped point cloud is voxelized to obtain a cleaner look with a uniform density. Additional static rendered images are also shown in Figures 4 and 5.

5. Discussion

In this section, we discuss the challenges we encountered that justify the sensors being employed and the proposed methodology. We also outline the main limitations of our PoC and discuss directions for future work.

Challenges The generation of this digital twin highlighted several surgery-specific challenges. Firstly, the frequent use of glass and metal surfaces in operating rooms poses a significant challenge for optical sensors and systems due to their reflectivity. We tested two commercial 3D scanners and a photogrammetry approach for the reconstruction of the operating room, but all failed to reconstruct the glass-covered walls or the metal operating tables. These early results also motivated us to use a laser scanner for the generation of a reference point cloud, which greatly simplified the registration of all entities in a common coordinate frame.

Secondly, the 3D reconstruction of human anatomy requires high-resolution imaging to capture its complex geometry and fine detailed texture. The resolution of the *Faro* laser, designed to capture large objects and environments, has shown to not be sufficient. This motivates the use of photogrammetry to reconstruct the operating table and anatomy, a technology that shows a good compromise between acquisition time and reconstruction accuracy at this scene scale.

Thirdly, the surgical scrubs posed a challenge for human body pose estimation method and specifically for the key-point detector, which were trained on humans wearing casual clothing. Refining the pretrained models on medical staff wearing scrubs, e.g. using the MVOR dataset [45], could yield a domain-specific model with an improved performance.

Limitations Our current implementation has several limitations that need to be addressed. Firstly, the photogrammetry-based reconstruction of the incision does not capture any dynamics, so the reconstruction of a full surgery would require multiple captures at different key steps as a minimum. The time-consuming capture process makes this approach unfeasible for real surgery. Instead, recent dynamic 3D reconstruction approaches based on NeRFs [38] or Gaussian splatting [47] could be utilized to reconstruct dynamic surfaces in the scene and specifically the incision. The limited changes of the patient anatomy during the step of pedicle screw placement motivated us to rely on a photogrammetry approach for this prototype.

Secondly, our proof-of-concept assumes rigidity of the instrument table, operating table, and OR lamps and displays. Articulations and movements of the operating tables could instead be tracked by continuously registering the 3D model to the dynamic point cloud, via ML-based pose estimation methods [24], or by utilizing marker-based tracking. In a similar fashion, a CT-based 3D model of the patient anatomy could be registered to the spatio-temporal scene by either point cloud-based registration [26] or via marker-based tracking. Also, the static display contents in our PoC could be replaced by a screen recording or a lightweight state-based representation of the shown contents.

Thirdly, the available data streams are added independently to our digital twin. As a result, calibration errors and noise can cause inconsistencies in the shared spatio-temporal representation, e.g. a mismatch in the hand and instrument pose. These inconsistencies could be reduced by integrating sensors jointly or based on a learnt model, such that the spatio-temporal consistency and plausibility of the digital twin can be enforced.

Finally, the generation of this prototype was time-consuming due to the lack of automated processes. Several steps of our pipeline, such as the registration of overhead devices, the capture of close-range photographs for the photogrammetric reconstruction of the operating table, or the registration of the anatomic model were conducted by hand. To enable a systematic and efficient generation of SDT, these manual steps need to be automated.

Future work Evolving our presented PoC into a complete SDT requires the integration of further sensors, an automated interpretation to generate semantic labels, and ideally the inclusion of prior knowledge. First, the integration of additional sensors like microphones, patient vitals, and medical imaging ensure that the digital twin accurately captures the available information at a time. Naturally, the most relevant sensors to monitor the patient already exist in today's OR, which reduces the cost of integrating additional sensors. Second, the data streams from all sensors need to be automatically analysed and interpreted to extract semantics. The method of estimating the surgeon's

body pose is representative for further extensions of our digital twin with estimated semantic annotations, for example from segmentation [31], surgical scene graphs [27], anatomical landmarks [51], or surgical phase detection [7]. Last, the integration of prior knowledge in form of physical, (bio-)mechanical [8], or behavioral models [50] is needed to extend the presented spatio-temporal reconstruction to a comprehensive digital twin.

6. Conclusion

In this work, we presented a proof-of-concept for surgery digitalization. We outlined the potential of SDTs for education, training data generation, simulation and closed-loop optimization, and automation of surgical tasks such as planning and reporting. We proposed a methodology to obtain an SDT encompassing the most relevant entities in surgery.

In contrast to related works that manually craft a virtual environment to simulate surgeries, our approach focuses on the capture of a real surgery. In its current state, our prototype can already be used to capture and re-render surgical steps or simple interventions for educational purposes, for example in the form of training videos or interactive virtual reality (VR)-based applications.

Our PoC is a step towards the systematic capture of surgeries, which may be used to collect a large dataset of digitized surgeries, including rare pathological cases and other infrequent events such as unforeseen complications or surgical errors. Moreover, the generated SDTs can provide a realistic environment for the training of ML-based models and robotic agents with a reduced sim-to-real gap. In the long run, holistic approaches to surgery digitization may boost the performance of state-of-the-art methods in computer-assisted surgery due to comprehensive representations of the current state of the surgery. We hope that our work motivates further research on automated methods for surgery digitization and the creation of SDTs.

Data availability

The data will be made available on our project page <https://jonashein.github.io/surgerydigitization/>.

Acknowledgements

This work has been supported by the OR-X - a swiss national research infrastructure for translational surgery - and associated funding by the University of Zurich and University Hospital Balgrist, as well as by the InnoSuisse PROFICIENCY grant. The study was conducted according to the guidelines of the Declaration of Helsinki, and approved by the local ethical committee (KEK Zurich BASEC No. 2021-01196). We thank Jan David Grunder for his support in the creation of the CAD model of the OR.

References

- [1] Petteri Aimonen. Fast and easy focus stacking, 2022. 5
- [2] Yotam Barnoy, Molly O'Brien, Will Wang, and Gregory D. Hager. Robotic surgery with lean reinforcement learning. *CoRR*, abs/2105.01006, 2021. 2
- [3] Marco A Barreto, Jorge Perez-Gonzalez, Hugh M Herr, and Joel C Huegel. Aracam: A rgb-d multi-view photogrammetry system for lower limb 3d reconstruction applications. *Sensors*, 22(7):2443, 2022. 3
- [4] Øystein Bjelland, Bismi Rasheed, Hans Georg Schaathun, Morten Dinhoff Pedersen, Martin Steinert, Alf Inge Hellevik, and Robin T. Bye. Toward a digital twin for arthroscopic knee surgery: A systematic review. *IEEE Access*, 10:45029–45052, 2022. 3
- [5] Guido Caccianiga, Julian Nubert, Marco Hutter, and Katherine J Kuchenbecker. Dense 3d reconstruction through lidar: A comparative study on ex-vivo porcine tissue. *arXiv preprint arXiv:2401.10709*, 2024. 3
- [6] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 4, 6
- [7] Tobias Czempiel, Magdalini Paschali, Daniel Ostler, Seong Tae Kim, Benjamin Busam, and Nassir Navab. Opera: Attention-regularized transformers for surgical phase recognition. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2021*, pages 604–614. Springer, 2021. 8
- [8] Tien Tuan Dao, Philippe Pouletaut, Áron Lazáry, and Marie Christine Ho Ba Tho. Multimodal medical imaging fusion for patient specific musculoskeletal modeling of the lumbar spine system in functional posture. *Journal of Medical and Biological Engineering*, 37:739–749, 2017. 8
- [9] Shounak Datta, Yanjun Li, Matthew M. Ruppert, Yuanfang Ren, Benjamin Shickel, Tezcan Ozrazgat-Baslantı, Parisa Rashidi, and Azra Bihorac. Reinforcement learning in surgery. *Surgery*, 170(1):329–332, 2021. 2
- [10] Mitchell Doughty and Nilesh R Ghugre. Hmd-egopose: Head-mounted display-based egocentric marker-less tool and hand pose estimation for augmented surgical guidance. *International Journal of Computer Assisted Radiology and Surgery*, 17(12):2253–2262, 2022. 3
- [11] Philip Edgcumbe, Philip Pratt, Guang-Zhong Yang, Christopher Nguan, and Robert Rohling. Pico lantern: Surface reconstruction and augmented reality in laparoscopic surgery using a pick-up laser projector. *Medical image analysis*, 25(1):95–102, 2015. 3
- [12] Hubertus Feußner and Adrian Park. Surgery 4.0: the natural culmination of the industrial revolution? *Innovative Surgical Sciences*, 2(3):105–108, 2017. 2
- [13] Beerend GA Gerats, Jelmer M Wolterink, and Ivo AMJ Broeders. Dynamic depth-supervised nerf for multi-view rgb-d operating room videos. In *International Workshop on Predictive Intelligence In Medicine*, pages 218–230. Springer, 2023. 3
- [14] Michael Grieves. Digital twin: manufacturing excellence through virtual factory replication. *White paper*, 1(2014): 1–7, 2014. 1, 3
- [15] Andrew A. Gumbs, Vincent Grasso, Nicolas Bourdel, Roland Croner, Gaya Spolverato, Isabella Frigerio, Alfredo Illanes, Mohammad Abu Hilal, Adrian Park, and Eyad Elyan. The advances in computer vision that are enabling more autonomous actions in surgery: A systematic review of the literature. *Sensors*, 22(13):4918, 2022. 2
- [16] Gregory D. Hager, Lena Maier-Hein, and S. Swaroop Vedula. Chapter 38 - Surgical data science. In *Handbook of Medical Image Computing and Computer Assisted Intervention*, pages 931–952. Academic Press, 2020. 3
- [17] Jonas Hein, Nicola Cavalcanti, Daniel Suter, Lukas Zingg, Fabio Carrillo, Lilian Calvet, Mazda Farshad, Marc Pollefeys, Nassir Navab, and Philipp Fünfstahl. Next-generation surgical navigation: Marker-less multi-view 6dof pose estimation of surgical instruments, 2023. 3, 5, 6
- [18] Felix Holm, Ghazal Ghazaei, Tobias Czempiel, Ege Özsoy, Stefan Saur, and Nassir Navab. Dynamic scene graph representation for surgical video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 81–87, 2023. 2, 3
- [19] Sascha Jecklin, Carla Jancik, Mazda Farshad, Philipp Fünfstahl, and Hooman Esfandiari. X23D - intraoperative 3d lumbar spine shape reconstruction based on sparse multi-view x-ray data. *J. Imaging*, 8(10):271, 2022. 2
- [20] Abdolrahim Kadkhodamohammadi, Afshin Gangi, Michel de Mathelin, and Nicolas Padoy. Articulated clinician detection using 3d pictorial structures on RGB-D data. *Medical Image Anal.*, 35:215–224, 2017. 2
- [21] Abdolrahim Kadkhodamohammadi, Afshin Gangi, Michel de Mathelin, and Nicolas Padoy. A multi-view rgb-d approach for human pose estimation in operating rooms. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 363–372. IEEE, 2017. 2
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 3
- [23] Heikki Laaki, Yoan Miche, and Kari Tammi. Prototyping a digital twin for real time remote control over mobile networks: Application of remote surgery. *Ieee Access*, 7:20325–20336, 2019. 2
- [24] Xiaolong Li, He Wang, Li Yi, Leonidas J Guibas, A Lynn Abbott, and Shuran Song. Category-level articulated object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3706–3715, 2020. 8
- [25] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *CVPR*, 2023. 3
- [26] Florentin Liebmann, Marco von Atzigen, Dominik Stütz, Julian Wolf, Lukas Zingg, Daniel Suter, Nicola A Cavalcanti, Laura Leoty, Hooman Esfandiari, Jess G Snedeker, et al. Automatic registration with continuous pose updates for marker-less surgical navigation in spine surgery. *Medical Image Analysis*, 91:103027, 2024. 8

- [27] Chen Lin, Zhenfeng Zhu, Yawei Zhao, Ying Zhang, Kunlun He, and Yao Zhao. Sgt++: Improved scene graph-guided transformer for surgical report generation. *IEEE Transactions on Medical Imaging*, 2023. **2, 8**
- [28] Yifan Liu, Chenxin Li, Chen Yang, and Yixuan Yuan. Endogaussian: Gaussian splatting for deformable surgical scene reconstruction. *arXiv preprint arXiv:2401.12561*, 2024. **3**
- [29] Hannah Lonsdale, Geoffrey M Gray, Luis M Ahumada, Hannah M Yates, Anna Varughese, and Mohamed A Rehman. The perioperative human digital twin. *Anesthesia & Analgesia*, 134(4):885–892, 2022. **3**
- [30] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. **4, 6**
- [31] Jun Ma, Yuting He, Feifei Li, Lin Han, Chenyu You, and Bo Wang. Segment anything in medical images. *Nature Communications*, 15(1):654, 2024. **8**
- [32] Lena Maier-Hein, Swaroop S Vedula, Stefanie Speidel, Nassir Navab, Ron Kikinis, Adrian Park, Matthias Eisenmann, Hubertus Feussner, Germain Forestier, Stamatia Giannarou, et al. Surgical data science for next-generation interventions. *Nature Biomedical Engineering*, 1(9):691–696, 2017. **3**
- [33] Lena Maier-Hein, Matthias Eisenmann, Duygu Sarikaya, Keno März, Toby Collins, Anand Malpani, Johannes Fallert, Hubertus Feussner, Stamatia Giannarou, Pietro Mascagni, et al. Surgical data science—from concepts toward clinical translation. *Medical image analysis*, 76:102306, 2022. **3**
- [34] Randi Q. Mao, Lucy Lan, Jeffrey Kay, Ryan Lohre, Olufemi R. Ayeni, Danny P. Goel, and Darren de SA. Immersive Virtual Reality for Surgical Training: A Systematic Review. *Journal of Surgical Research*, 268:40–58, 2021. **2**
- [35] Pietro Mascagni, Deepak Alapatt, Luca Sestini, Maria S. Altieri, Amin Madani, Yusuke Watanabe, Adnan Alseidi, Jay A. Redan, Sergio Alfieri, Guido Costamagna, Ivo Boskoski, Nicolas Padoy, and Daniel A. Hashimoto. Computer vision in surgery: from potential to clinical value. *npj Digit. Medicine*, 5, 2022. **2**
- [36] Xavier Maurice, Chadi Albitar, Christophe Doignon, and Michel de Mathelin. A structured light-based laparoscope with real-time organs’ surface reconstruction for minimally invasive surgery. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 5769–5772. IEEE, 2012. **3**
- [37] Rory McCloy and Robert Stone. Virtual reality in surgery. *Bmj*, 323(7318):912–915, 2001. **1**
- [38] Marko Mihajlovic, Sergey Prokudin, Marc Pollefeys, and Siyu Tang. Resfields: Residual neural fields for spatiotemporal signals. *arXiv preprint arXiv:2309.03160*, 2023. **8**
- [39] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. **3**
- [40] Thomas Neumuth. Surgical process modeling. *Innovative surgical sciences*, 2(3):123–137, 2017. **3**
- [41] Ege Özsoy, Evin Pınar Örnek, Ulrich Eck, Federico Tombari, and Nassir Navab. Multimodal semantic scene graphs for holistic modeling of surgical procedures. *arXiv preprint arXiv:2106.15309*, 2021. **2, 3**
- [42] Mariam M. Raza, Kaushik P. Venkatesh, James A. Diao, and Joseph C. Kvedar. Defining digital surgery for the future. *npj Digit. Medicine*, 5, 2022. **2**
- [43] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), 2017. **4, 6**
- [44] Hongchao Shu, Ruixing Liang, Zhaoshuo Li, Anna Goodridge, Xiangyu Zhang, Hao Ding, Nimesh Nagururu, Manish Sahu, Francis X Creighton, Russell H Taylor, et al. Twin-s: a digital twin for skull base surgery. *International Journal of Computer Assisted Radiology and Surgery*, pages 1–8, 2023. **3**
- [45] Vinkle Srivastav, Thibaut Issenhuth, Kadkhodamohammadi Abdolrahim, Michel de Mathelin, Afshin Gangi, and Nicolas Padoy. Mvor: A multi-view rgb-d operating room dataset for 2d and 3d human pose estimation. 2018. **7**
- [46] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3295–3306, 2023. **3**
- [47] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Wang Xinggang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023. **8**
- [48] Jiaqi Xu, Bin Li, Bo Lu, Yun-Hui Liu, Qi Dou, and Pheng-Ann Heng. Surrol: An open-source reinforcement learning centered and dvrk compatible platform for surgical robot learning. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1821–1828. IEEE, 2021. **1**
- [49] Kun Yuan, Manasi Kattel, Joel L Lavanchy, Nassir Navab, Vinkle Srivastava, and Nicolas Padoy. Advancing surgical vqa with scene graph knowledge. *arXiv preprint arXiv:2312.10251*, 2023. **3**
- [50] Kaifeng Zhao, Yan Zhang, Shaofei Wang, Thabo Beeler, and Siyu Tang. Synthesizing diverse human motions in 3d indoor scenes. In *ICCV*, pages 14692–14703, 2023. **2, 8**
- [51] Heqin Zhu, Qingsong Yao, Li Xiao, and S Kevin Zhou. You only learn once: Universal anatomical landmark detection. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24*, pages 85–95. Springer, 2021. **8**
- [52] Ege Özsoy, Tobias Czempel, Felix Holm, Chantal Pellegrini, and Nassir Navab. LABRAD-OR: Lightweight Memory Scene Graphs for Accurate Bimodal Reasoning in Dynamic Operating Rooms. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, pages 302–311, Cham, 2023. Springer Nature Switzerland. **3**