The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Surgical Case Identification for an Image-Guided Interventional System

Tamás Haidegger, Peter Kazanzides, Balázs Benyó, Levente Kovács and Zoltán Benyó

Abstract— **Image-guided surgery offers great advantages to surgeons through the possibility to track tools in 3D space and to navigate based on the virtual model of the patient. In the case of robot-assisted procedures, both the inherent accuracy of the system components and the quality of the registration procedures are critical to provide high precision treatment delivery. One of the major barriers towards more technology-integrated procedures is the fact that alterations in the operating room environment can fundamentally change the performance of the system, decrease the accuracy, and therefore pose significant danger to the patient. Surgical events from the control point of view may include motion of the robot, motion of the camera, or motion of the patient. The paper describes a new concept to treat these events, to track and automatically compensate for abrupt changes that may affect the accuracy of a robot-integrated interventional system. Our solution is to use all available information at a given time, including the intra-operative tracker's internal base frame, to distinguish between different surgical events. The concept has been developed and tested on the neurosurgical robot system at the Johns Hopkins University. Initial experiments performed on data recordings from simulated scenarios showed that the algorithm was able to correctly identify the cases.**

I. INTRODUCTION

Computer-Integrated Surgery (CIS) refers to theory and technology intended to improve the efficiency and accuracy of health care delivery. CIS means the combination of innovative algorithms, robotic devices, imaging systems, sensors and human–machine interfaces to work cooperatively with physicians in the planning and execution of surgical procedures [1]. A subfield of CIS is called Image-Guided Surgery (IGS), where the digital system is not necessarily involved in the physical part of the operation, but improves the quality of surgery through better visualization or guidance. IGS means the accurate registration (correlation and mapping) of the pre-operative (MR, CT) or intra-operative (ultrasound, fluoroscopy) data set of the patient to the actual location within the operating room (OR). It may be useful to register further devices, e.g., robots to the setup, providing not only free-hand navigation, but also accurate positioning, or guidance of a tool.

T. Haidegger, B. Benyó, L. Kovács and Z. Benyó are with Dept. of Control Engineering and Information Technology, Budapest University of Technology and Economics, Magyar tudósok krt. 2, Budapest, Hungary haidegger@, bbenyo@, lkovacs@, benyo@iit.bme.hu

Peter Kazanzides is with the Johns Hopkins University, Center for Computer Integrated Surgical Systems and Technology, (CISST ERC), N. Charles 3400, Baltimore, MD, USA pkaz@jhu.edu

This work was supported by the National Office for Research and Technology (NKTH), Hungarian National Scientific Research Foundation grant OTKA CK80316.

A. Imperfections of CIS systems

Robots can support surgeons with advanced targeting, steady positioning and precise task execution, surpassing human capabilities. Effectiveness—accuracy of treatment delivery—of an interventional system is crucial, especially when it is operating semi-autonomously. Researchers are focusing on improving systems through eliminating errors that may originate in:

- CT/MR imaging errors and resolution issues
- volume model generation errors
- inaccurate treatment planning
- errors introduced by hardware fixation
- noise on intra-operatively collected data
- registration errors
- inherent inaccuracies of surgical tools and actions
- patient motion

The overall application accuracy of a system can be affected in a complex manner by all of these. Errors can vary in frequency, amplitude and correlation, therefore combined hardware and software solutions are required to effectively reduce their effects. Probably the most dangerous deviations are the small, incremental changes in the setup and the OR environment, causing errors that are hard to detect immediately by the human operator or assistant.

B. State of the art in patient motion tracking

The event of patient motion occurs when the body moves relative to the base of the equipment executing the surgical plan. The fundamental problem with patient motion is that without proper identification and compensation, the whole surgical plan may become obsolete. If noticed in time, re-registration is recommended to avoid damaging the patient. However, re-registration is usually time consuming and cumbersome, therefore it should be avoided, whenever possible. From the technical point of view, many sources of errors can be represented as patient motion. The main sources of external (i.e., excluding physiological) patient motion during surgery include:

- large forces applied by surgeon (e.g., bone milling)
- bumping into the operating table
- leaning against the patient
- inadequate fixation
- equipment failure

The robot's position information and the tracking data must be kept consistent throughout the operation. Currently deployed systems use different approaches, but most of them relies of rigid anchoring. Smaller robots, such as the Smart-Assist [2] (Mazor Surgical Technologies Inc., Caesarea, Israel) or the MBARS [3] may be bone-mounted. This requires more invasive fixation on the patient side (bone screws), and large forces may still cause relative motion between the patient and the tool. In orthopedics, there are significant interactions, making it necessary to use more invasive fixations. Using a large robot with strong, rigid attachments to the patient may introduce the danger of serious tissue damage. The ROBODOC system [4] (Curexo Tech. Inc., Fremont, CA) was the first Food and Drug Administration (FDA) approved automated bone milling robot for hip replacement, and it uses bone-attached fixation together with a bone motion sensor to detect fixation failures. If the bone moves more than 2 mm despite the fixation, the system halts, and calls for re-registration. There is a clear trend in surgical applications to shift towards less invasive solutions.

One option is to use multiple Dynamic Reference Bases (DRB—rigid bodies constructed from trackable markers), to follow the motion of the robot base and the patient separately. Unfortunately, not every tracking system supports this, and it may cause difficulties to maintain the line-of-sight without disturbing the physician. Extending the active workspace of a tracking system may result in higher inherent accuracies due to the inhomogeneity of their fields. Some commercial systems combine surface-mounted and in-body fiducials to track external and physiological organ motion, though it requires a separate operation just to place the markers. A successful example is the CyberKnife radiation therapy system (Accuray Inc., Sunnyvale, CA) that can track both skin motion through a special suit and organ motion by taking bi-plane x-ray images and locating fiducials (gold beads in this case) that were implanted preoperatively [5].

Robotic setups could include accelerometers and gyroscopes to detect sudden changes; however there is need for electronic coupling and the resolution may not be sufficient for proper compensation. Besides, these would increase the costs and complexity of the system. CCD cameras can survey the OR, and image processing techniques could solve the localization problem, but the resolution may not be high enough, and it might have significant hardware requirements.

Further, dynamic registration and correction for patient motion has been implemented with PET scans [6], [7] to improve image quality through compensated reconstruction. However, these setups only considered rigid environment, where neither the camera, nor the PET gantry moved.

While significant effort has been invested to describe the surgical workflow with mathematical models [8], relatively few projects have dealt with the modeling of the OR setup and environment in general. Dynamic correction for unforseen events remains a significant challenge with currently used typical intra-operative navigation systems.

II. GENERAL CONCEPT FOR MOTION COMPENSATION

A tracking-based minimally invasive concept for patient motion compensation was proposed earlier [9] to support systems with less rigid fixation setup or limited surgical navigation capabilities. It is based on the principle that during regular operation, the position of the surgical tool mounted

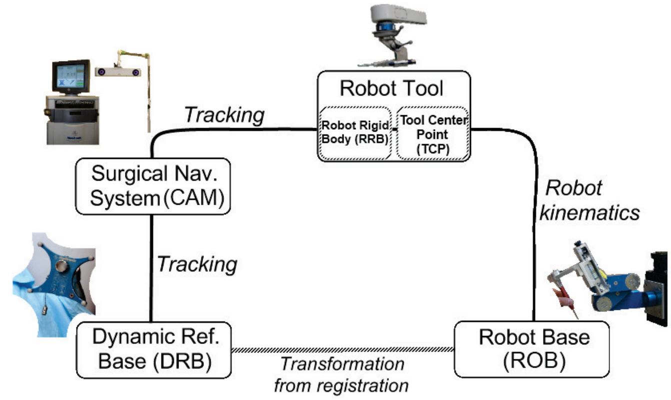

Fig. 1. General control concept of image-guided robotic systems

on an IG robot can be controlled precisely, once its location is known relative to the base coordinate system. The base frame can be chosen arbitrarily as long as it is connected to the robot frame (and possible other control frames, such as the navigation systems) through known homogenous transformations. A generic robot-integrated image-guided system's schematic diagram is shown in Fig. 1. The nodes represent control frames, and the lines homogenous transformations. The camera system is able to track in 3D both the motion of the patient and the tip of the robot. Typically, the origin of the trackable marker (Robot Rigid Body—RRB) and the Tool Center Point (TCP) of the robot might be different. In this case, the fixed transformation connecting them is obtained from calibration, if not known a priori.

Control signals are practically computed in the DRB frame (patient's coordinate system), and are transformed to the robot base frame (ROB) executing a series of homogenous transformations through the camera's coordinate frame (CAM):

$$\text{Control}|_{\text{ROB}} = {}^{\text{DRB}}_{\text{CAM}}T \cdot {}^{\text{CAM}}_{\text{RRB}}T \cdot {}^{\text{RRB}}_{\text{TCP}}T \cdot {}^{\text{TCP}}_{\text{ROB}}T \cdot \text{Control}|_{\text{DRB}}. \quad (1)$$

It is possible to close the entire control loop in Fig. 1 through calibration, acquiring the transformation between the ROB and the DRB frames. This can be computed under stationary conditions (during the setup) by calculating:

$${}^{\text{DRB}}_{\text{ROB}}T = {}^{\text{DRB}}_{\text{CAM}}T \cdot {}^{\text{CAM}}_{\text{RRB}}T \cdot {}^{\text{RRB}}_{\text{TCP}}T \cdot {}^{\text{TCP}}_{\text{ROB}}T. \quad (2)$$

Throughout the surgery, when unintentional motions of the patient with respect to the robot are detected—i.e., deviations from the original ROB to DRB transformation—it is possible to compensate for the change by recomputing (2). This ensures that the original surgical plan stays valid.

The effectiveness of the concept depends on the inherent accuracy of the different elements—the robot and the navigation system—and the registration. First and foremost, the spatial localization can be corrupted by noise. In the case of the most typically used optical trackers, noise varies with the used marker type, the lighting conditions and the position and angle of the rigid bodies in the cameras' field of view. In the case of electromagnetic tracking, susceptibility to ferromagnetic materials in the proximity of the sensor

TABLE I

DIFFERENT SURGICAL MOTION SCENARIOS TOGETHER WITH THE REQUIRED ACTION TO TAKE IN THAT DISTINCT CASE.

	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6	Case 7	Case 8
Camera Moving (Y/N)	Y	Y	Y	Y	N	N	N	N
Patient Moving (Y/N)	Y	Y	N	N	Y	Y	N	N
Robot Moving (Y/N)	Y	N	Y	N	Y	N	Y	N
Desired action	Re-register	Slow robot	Slow robot	Slow robot	Slow robot	Slow robot		

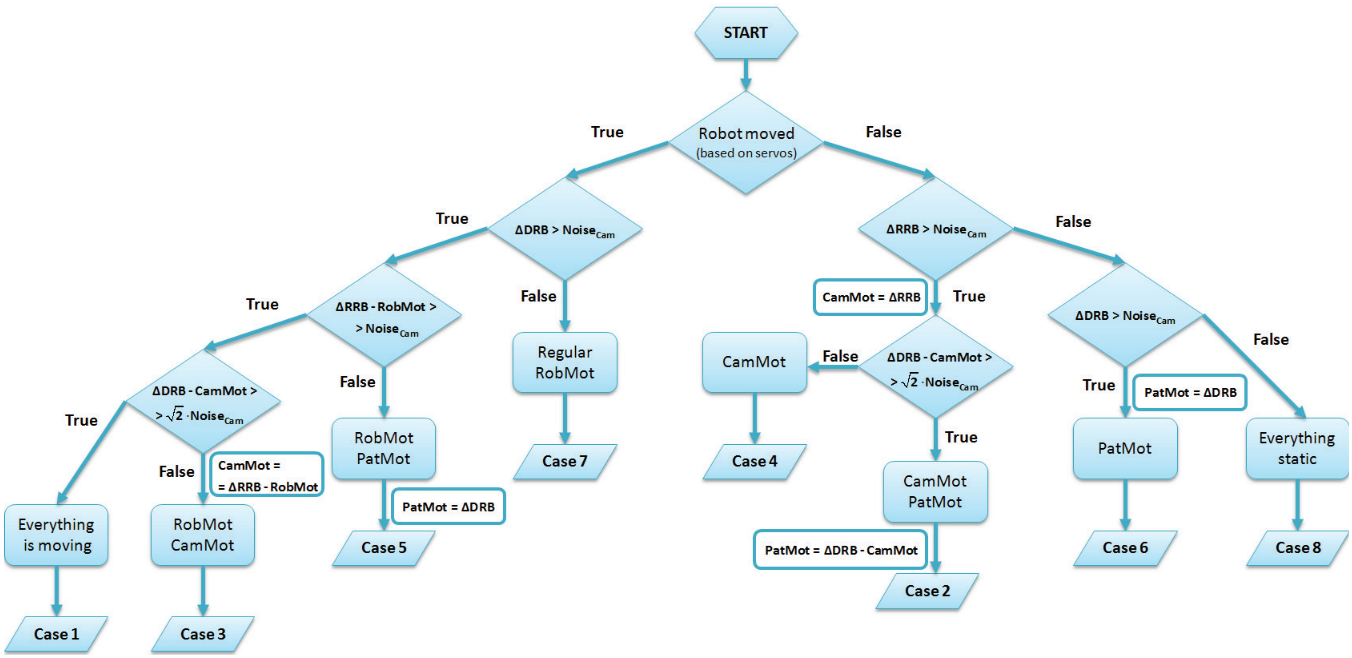

Fig. 2. The decision-making chart showing the evaluation of different OR situations.

can cause significant distortions. Beyond measurement noise, latency is also a major problem, making it harder to close the control loop for compensation through (2). Furthermore, robotic systems typically run in 20–100 Hz cycles at the highest control level; however, commercially available tracking systems are not capable of more than 5–60 Hz data acquisition, depending on their modality. Slower sampling rate only allows less frequent and therefore less accurate localization.

A. New Surgical Case Identification Concept

In an integrated surgical robotic setup (Fig. 1) changes in the operating room setup are due to the different elements moving relative to each other. Alterations can be modeled as distinct events of camera motion (CamMot), patient motion (PatMot), the regular motion of the robot (RobMot), or the arbitrary combination of these. It can be assumed that the robot is anchored rigidly to the ground, therefore its base cannot move during the procedure. The deriving eight categories—surgical cases from the control point of view—are listed in Table I. From the application point of view, these events can be categorized either as *free motion* (regular safe operation), or *potentially dangerous situation* (when something is moving, but we are still able to compensate for it) and finally *forbidden event*, when re-registration is necessary to regain accuracy in control. During potentially dangerous events, it is required to slow the robot proportionally to the factor of the risk, to allow for the collection of multiple measurements to reduce the probability of damage. This generic description allows for a unified handling of any IG system complying with the structure outlined in Fig. 1.

A robust method is required to identify the different surgical cases, and an optimal position estimation technique for each can be derived separately. According to Table I, Case 8 refers to the completely stationary state, and during regular operation only the robot is moving and performing the desired tasks (Case 7). To support flexibility and line of sight, time to time, the camera might be moved or relocated. General practice allows for the free adjustment of the camera position, as only the ${}^{\mathrm{DRB}}_{\mathrm{CAM}}T^{-1} \cdot {}^{\mathrm{RRB}}_{\mathrm{CAM}}T$ is used to generate control signals. It is crucial to reliably separate Case 1, 2, 5 and 6 from the rest, as patient motion is the critical event we wish to compensate for.

To apply the concept described above, the position changes of the DRB (i.e., the amount of patient motion) must be known. To be able to determine it, all available information of the system should be used. The novelty of our concept is that it incorporates in the calculations the navigation system's inner coordinate base frame as well. To rely on the camera internal frame (CAM), it must be assured that the camera itself does not move during the evaluation period. This option has not been investigated by other groups, as we know. In a general OR setup, this may be inconvenient to achieve through the rigid fixation of the camera system, therefore dynamic event recognition is required. It is important to note that camera repositioning occurs rarely during a procedure, although it must be monitored and compensated accurately. Under the given condition, Cases 5, 6, 7 and 8 can be best handled by using the camera frame as an individual reference for patient motion (basically equivalent to having an additional DRB). This allows for the easy handling of Cases 5 and 6 using CAM as a reference, and assuming that we have an accurate recording of the robot's motion. When the robot is static, we can use its position as a temporary reference frame to determine possible CamMot (Cases 2 and 4). Subtracting the camera motion from the recorded difference in the DRB's position, PatMot can be identified and numerically determined through $\Delta({}^{\text{CAM}}_{\text{DRB}}T^{-1} \cdot {}^{\text{CAM}}_{\text{RRB}}T) = 0+$ PatMot, where Δ means the increment between two control cycles. If the robot is moving, the equation is extended with the precisely known displacement of the robot's tip in the last cycle (RobMot)—transformed to the CAM frame (Case 3).

B. Specific Issues With Motion Compensation

In practice, the aforementioned errors complicate the situation. Each decision point becomes less reliable due to the distribution of random errors. Fig. 2 summarizes the new surgical event decision concept, including the error margins. Every time the camera measurement is involved in the decision condition, an error tolerance zone should be considered—proportional to the noise parameters (Fig. 3). In accordance with the literature, we consider unisotropic Gaussian noise distribution [10]. The $\text{Noise}_{\text{Cam}}$ parameter should be determined for every device based on manufacturer's specification or experiments. Typically, intra-operative tracking systems provide sub-millimeter accuracy even with the consistent calculation of ${}^{\text{DRB}}_{\text{CAM}}T^{-1} \cdot {}^{\text{RRB}}_{\text{CAM}}T$. If the sum of two independent measurements are used from the camera (e.g., to distinguish Case 2 and Case 4), the resulting noise distribution is the convolution of the component noise distributions. (If X_i has Gaussian distribution with parameters $N(0, \sigma^2)$, then $\sum_{i=1}^{n} X_i$ also has Gaussian distribution with parameters $N(0, n\sigma^2)$.)

There are certain assumptions we make in order to treat the system in a unified way:

- information about the state of the robot is available
- the robot is more precise than the navigation system
- the noise parameters of the components are known
- the latencies between the system components are known

In general, we use all relevant information given at a point of measurement to make optimal decision on what surgical conditions apply to the current OR setup. First, we compensate for average latencies occurring in the system. Next, we need to distinguish whether the robot is moving or not. Usually, this information is directly available through the servo feedback. Averaging over time may help reducing the error, however, the safe control of the system during patient motion events must also be addressed. Alteration might either be a sudden, or a gradual event (drift or shift), and it can take seconds for the DRB to settle. It must be ensured that in the mean time the robot does not move beyond a user-defined safety margin—typically 1 mm—it can be scaled down accordingly. To dampen the effect of noise (primarily in the camera system), filtering can be applied to the signal. Due to the sudden changes (relative to the time resolution of a navigation system), we only perform average filtering while the system is static. Based on the average noise distribution of a camera system, we defined the window size for averaging. Position and orientation values of the rigid bodies are averaged through a maximum of 50 cycles. If patient motion occurs, these values (averaged over avg_old cycles) will be stored, and motion event will be computed relative to these. When the system stabilizes, a new average is cumulated in every cycle (avg_new). When $avg\_new == avg\_old$, we replace the older value with the new one. Fig. 3 illustrates a single patient motion event to support this concept (in 1D for clarity). While performing PatMot and CamMot computations based on Fig. 2, the following values are stored:

$$\Delta\text{RRB} = \Delta\,{}^{\text{RRB}}_{\text{CAM}}T = {}^{\text{RRB}}_{\text{CAM}}T|_{avg\_new} - {}^{\text{RRB}}_{\text{CAM}}T|_{avg\_old} \quad (3)$$

$$\Delta\text{DRB} = \Delta\,{}^{\text{DRB}}_{\text{CAM}}T = {}^{\text{DRB}}_{\text{CAM}}T|_{avg\_new} - {}^{\text{DRB}}_{\text{CAM}}T|_{avg\_old.} \quad (4)$$

Based on these, we can update the ${}^{\text{DRB}}_{\text{ROB}}T$ and ${}^{\text{CAM}}_{\text{ROB}}T$ transformations, originally determined through the initial robot registration phase.

In the unlikely event when everything is moving at the same time, the algorithm should still be able to determine the individual positions, but the overall uncertainty of the estimation would be too large, therefore it is recommended to re-register the system. In currently deployed robots (such as the ROBODOC), this is a common practice to require re-registration if deviation from the original registration is noticed. In addition, if patient motion and camera motion occurs within 5 updates, re-registration is required again, due to the decreased accuracy of the shorter averaging period.

A simulation environment was developed under MATLAB R2008b (Mathworks Inc., Natick, MA) to verify the new concept. A setup similar to Fig. 1 was defined, where a robot and the patient are both tracked. Simulation parameters were chosen to be in the range of a general surgical robot regarding both range and speed of motion. Patient and camera motion events were determined based on previous real-life observations. The total length of the simulation was 160 s, assuming 10 Hz control cycle. RobMot, CamMot and PatMot events were permutated, leading to the following Case series in accordance with Table I: 8–7–3–4–6–5–2–1, each lasting for 20 s, and consisting of a larger scale motion and two smaller, but faster motion. First simulations were performed under idealized conditions, with no noise and zero

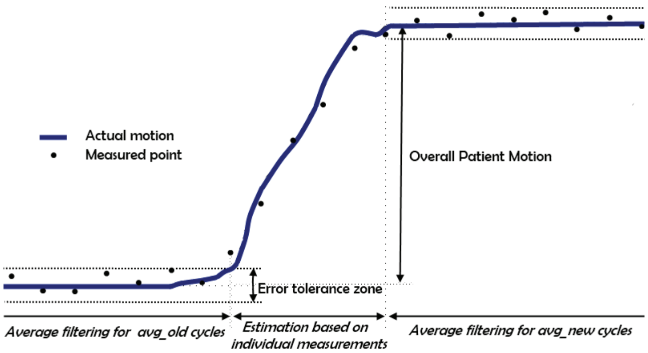

Fig. 3. Scheme of averaging for patient motion events

latency. The results showed the perfect identification of the surgical cases. The overall effectiveness is expressed with the percentage of correct Case identifications was 99.5%.

III. EVALUATION STUDY

A. The JHU Neurosurgical Robot

The concept introduced in the previous section was tested on a neurosurgical robot system developed at the Center for Computer-Integrated Surgical Systems and Technology (CISST ERC) at the Johns Hopkins University [11]. The project uses a modified robot and navigation making it capable of helping and increasing the performance of human surgeons (Fig. 4). The NeuroMate manipulator (formerly marketed by Integrated Surgical Systems Inc., Sacramento, CA) is a 5 degree-of-freedom (DOF) robot, originally built for stereotactic procedures. The StealthStation intra-operative navigation device is also commercially available (Medtronic Inc., Louisville, CO). It tracks the 3D position and orientation of sets of optical markers forming a rigid body. This version of the StealthStation only allows for the detection of two frames (i.e., a fixed reference frame and a moving probe). Both devices are approved by the FDA. An Anspach eMax 2 high-speed bone drilling surgical instrument (The Anspach Effort Inc., Palm Beach Gardens, FL) was attached to the tip of the robot through a 6 DOF force sensor (JR3 Inc., Woodland, CA, USA) to measure the forces and torques applied to the end-effector. The robot is guided in cooperative control mode for the removal of cranial bone on the skull base, while virtual fixtures (safety boundaries) are applied to protect critical anatomical structures [12]. The control software successfully integrates open source and proprietary components. It extensively relies on the cisst open source libraries (https://trac.lcsr.jhu.edu/cisst), developed for surgical robotic applications at Johns Hopkins.

B. Test Data for Evaluation

Unfortunately, there is very little published data on patient motion [13], [14]. For consistency, we examined several motion patterns correlating to the different events within the

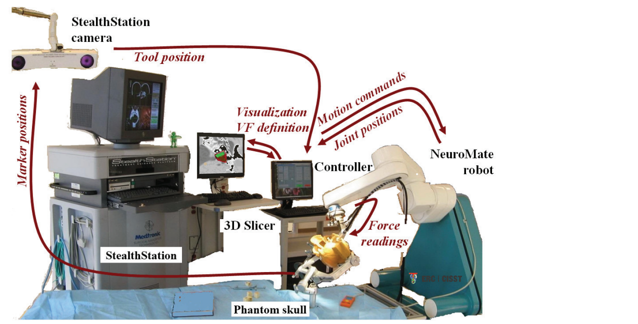

Fig. 4. Neurosurgical robot setup at the Johns Hopkins University, consisting of a manipulator, a surgical navigation system, a visualization console, and a robot controller computer.

OR, described in Section II. For evaluation purposes, we generated the following motion patterns:

- applying arbitrary forces on the cranial frame
- applying arbitrary forces on the robot
- applying arbitrary forces on the patient cart
- crisp motion of the DRB (bumping into it)
- drifting motion of the DRB (primarily rotation)

The initial experiments showed that the compliance of the cranial frame (Mayfield Infinity Skull Clamp) and the NeuroMate arm can already result in 0.85 mm and 1.55 mm RMS errors, respectively, therefore reducing the effectiveness of the treatment without proper compensation.

We performed test sequences according to the eight surgically relevant scenarios presented in Table I. These recordings were used for comparable evaluation of the new method. To be able to apply event based compensation, the specific parameters of the JHU system had been identified. We found that the inherent noise (standard deviation) of the Medtronic navigation system is $[0.0185, 0.0134, 0.0474]$ along x, y, z axes (in mm) and $[0.0029, 0.0009, 0.0029]$ in rotation along the same axes (in degrees), respectively. Spatial location within the workspace of the camera showed little deviation from these averaged values. On average, the StealthStation updates its interior position information in every 108 ms, and in total, the measurements arrive to the controller computer with 259.9 ms delay compared to the robot's signals.

C. Results and discussion

Preliminary results proved that our method managed to identify the most extreme events (Case 1 and 8) with 100% accuracy, and the more complex events with approximately 80% on average. Case 6 and 7 can be identified with a higher success rate, while Case 2 and 4 (and also Case 3 and 5) are harder to distinguish. The proposed gradually extending window averaging provides position estimation as soon as the dynamic reference frame (therefore the patient) stabilizes, and updates the registration transformations continuously, with increasing precision.

As the data sets were acquired on a physical system, imperfections of the different devices could lead to false clustering. The jitter in latency may also cause trouble, making it harder to distinguish surgical cases from each other.

Correlated motion patterns may further degrade accuracy. This may be an issue with actual bone milling with loose head fixation, when the robot's motion can induce patient motion, and the whole can look like a camera motion from the control point of view. Frequency domain identification of simultaneous motion types may help to separate Case 1 and 3. It should be mentioned that these cases happen seldom during regular clinical settings (i.e., the assistant should not adjust the camera while the surgeon is operating). The overall structure of the surgical case identification allows for a more generic, probability-based handling of the events that may lead to a more robust classification in the future. We plane to perform data collection on patient motions during actual surgeries to be able to better tune the parameters.

D. Limitations

The theoretical performance and robustness of the proposed method may be affected by many factors. While average communication latency can be compensated accurately, it evokes that overall prediction will be delayed with the maximum latency. This must be taken into consideration when designing the safety margins for the application. Sometimes, due to timing issues, latency may have a variance or jitter. Overall safety boundaries should be built in the surgical plans to tolerate maximum localization errors caused by occasional longer delays. If one component is significantly slower than the other, it may be hard to achieve real-time compensation. There is always a possibility to reduce the robot's speed to collect sufficient information, however, this may result in the unnecessary prolongation of the surgical procedure which is absolutely undesired. While false positive characterization of patient motion events may not mean inherent danger, it should be avoided to ensure smooth operation of the system, and the possibility to reduce operating time.

Both the inherent inaccuracies of the components and the registration procedures may introduce localization errors. In the case of an integrated system, this error can be magnified through the computation of a chain of homogenous transformations, where angular errors (in the estimation of rotation angles) will be multiplied out by the translational parameters. This should be handled through the extension of the method, involving probability based computation of the tool's location instead of using deterministic models.

IV. CONCLUSION

In the case of robot-assisted image-guided surgery or radiation therapy, precision is vital, therefore accurate positioning of the surgical tool is required. Different approaches have been experimented to robustly identify surgical events (such as patient motion), that may endanger the success of the operation or could cause damage to the patient. The aim of the research was to improve surgical robot navigation through decision-tree based event recognition, allowing to choose the best control option for the given case.

A new approach has been proposed for integrated IGS systems, using the intra-operative navigation device's internal coordinate base frame to better estimate the possible changes in the operating room environment. The advantage of the method is that it is arbitrary scalable to any improve different setups. Possible events have been categorized and defined as patient motion, camera motion, robot motion and the arbitrary combination of these. A measurement-based method was given to identify the actual case, and perform the desired action to enhance patient safety. The method was tested on an image-guided, cooperatively controlled neurosurgical robot developed at the Johns Hopkins University CISST ERC center. Different motion scenarios were simulated and recorded. The algorithm managed to robustly identify the most frequently occurring surgical cases.

V. ACKNOWLEDGMENTS

Special thanks to Sándor Győri for his insightful comments throughout the project.

REFERENCES

[1] R. Taylor and P. Kazanzides, "Medical Robotics and Computer-Integrated Interventional Medicine," *Advances in Computers: Emerging Technologies*, vol. 73, pp. 219–263, 2008.

[2] C. Plaskos, P. Cinquin, S. Lavallée, and A. Hodgson, "Praxiteles: a miniature bone-mounted robot for minimal access total knee arthroplasty." *The international journal of medical robotics and computer assisted surgery: MRCAS*, vol. 1, no. 4, pp. 67–79, 2005.

[3] A. Wolf, B. Jaramaz, B. Lisien, and A. DiGioia, "MBARS: mini bone-attached robotic system for joint arthroplasty." *The international journal of medical robotics and computer assisted surgery: MRCAS*, vol. 1, no. 2, pp. 101–121, January 2005.

[4] P. Kazanzides, "Robot assisted surgery: the ROBODOC experience," *Intl. Symposium on Robotics (ISR)*, vol. 30, pp. 281–286, 1999.

[5] K. Saito, M. Fujii, K. Kajiwara, and M. Suzuki, "Introducing sitetrack: continuous patient motion monitoring during stereotactic radiotherapy for the head." *Neurosurgery*, vol. 64, no. 2 Suppl, pp. 110–122, 2009.

[6] R. Fulton, S. Meikle, S. Eberl, J. Pfeiffer, C. Constable, and M. Fulham, "Correction for head movements in positron emission tomography using an optical motion-tracking system," *IEEE Transactions on Nuclear Science*, vol. 49, no. 4, pp. 116–123, 2002.

[7] A. Rahmim, O. Rousset, and H. Zaidi, "Strategies for motion tracking and correction in PET," *PET Clinics*, vol. 2, no. 2, pp. 251–266, 2007.

[8] U. Lahiri, R. F. Labadie, C. Liu, R. Balachandran, O. Majdani, and N. Sarkar, "A step toward identification of surgical actions in mastoidectomy." *IEEE Trans. on Bio-medical Engineering*, vol. 57, no. 2, pp. 479–87, February 2010.

[9] T. Haidegger, Z. Benyo, and P. Kazanzides, "Patient motion tracking in the presence of measurement errors." in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, Minneapolis, 2009, pp. 5563–5567.

[10] M. H. Moghari, B. Ma, and P. Abolmaesumi, "A theoretical comparison of different target registration error estimators." *MICCAI—International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 11, no. 2, pp. 1032–1040, January 2008.

[11] T. Xia, C. Baird, G. Jallo, K. Hayes, N. Nakajima, N. Hata, and P. Kazanzides, "An integrated system for planning, navigation and robotic assistance for skull base surgery," *International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 4, no. 4, pp. 321–330, 2008.

[12] L. Rosenberg, "Virtual fixtures: Perceptual tools for telerobotic manipulation," in *IEEE Virtual Reality Annual International Symposium*, 1993, pp. 76–82.

[13] B. Westermann and R. Hauser, "Online head motion tracking applied to the patient registration problem," *Journal of Image Guided Surgery*, vol. 5, no. 3, pp. 137–147, 2000.

[14] M. Ryan, R. Erickson, D. Levin, C. Pelizzari, R. MacDonald, and G. Dohrmann, "Frameless stereotaxy with real-time tracking of patient head movement and retrospective patientimage registration," *Journal of Neurosurgery*, vol. 85, pp. 287–292, 1996.