Constrained Visual-Inertial Localization With Application And Benchmark in Laparoscopic Surgery

Regine Hartwig[1,3], Daniel Ostler[1], Jean-Claude Rosenthal[2]
Hubertus Feußner[1], Dirk Wilhelm[1], Dirk Wollherr[3]

Abstract— We propose a novel method to tackle the visual-inertial localization problem for constrained camera movements. We use residuals from the different modalities to jointly optimize a global cost function. The residuals emerge from IMU measurements, stereoscopic feature points, and constraints on possible solutions in $SE(3)$. In settings where dynamic disturbances are frequent, the residuals reduce the complexity of the problem and make localization feasible. We verify the advantages of our method in a suitable medical use case and produce a dataset capturing a minimally invasive surgery in the abdomen. Our novel clinical dataset MITI is comparable to state-of-the-art evaluation datasets, contains calibration and synchronization and is available at [1].

I. INTRODUCTION

Many researchers have studied simultaneous localization and mapping (SLAM) in the computer vision, and robotics community [2], inspiring many innovative products. Applications such as virtual reality, augmented reality [3], unmanned vehicles, and autonomous robots [4] rely heavily on stable pose estimates based on SLAM systems.

Challenging situations include a dynamic environment which prevents the SLAM algorithms from being applicable in many use-cases. The works of [5–7] show that it is possible to add residuals specially designed for a use-case and jointly optimize a cost function. Motion priors, reprojection from a calibrated multi-camera setup, or additional measurements reduce the degrees of freedom on the solution space.

MIS has proven to have many medical advantages for the patient. Stable virtual/ augmented reality and assisting roboter arms could lead to enhanced hand-eye coordination in MIS and subsequently to a whole new way of conducting surgeries. Furthermore, SLAM is a game-changer for applications such as mapping detected metastasizing carcinoma for cancer staging, safety routines (e.g., collision avoidance), autonomous systems, image-guided surgery, input generation for artificial intelligence, and anatomic scene graphs. The huge disadvantages (magnetic distortion, line-of-sight, room consumption) disqualify available tracking equipment for most interventions. Data from this domain was acquired, and implementation shows the applicability of the proposed constrained visual-inertial SLAM method in this challenging dynamic surrounding.

[1] Authors are with Research Group MITI, Surgical Department, Technical University of Munich (TUM), Germany

[2] Authors are with Fraunhofer HHI, Berlin, Germany

[3] Authors are with Chair of Automatic Control Engineering, Department of Electrical and Computer Engineering, Technical University of Munich (TUM), Germany `regine.hartwig@tum.de`

The paper is structured as follows: Section IV presents the visual-inertial odometry algorithm with movement constraints. Section V discusses our novel dataset [1] and provides useful information for users. In Section VI we will discuss the implementation of the proposed algorithm, discuss properties of the calibrated data from our dataset, and evaluate the performance by experiments.

II. PROBLEM STATEMENT AND CONTRIBUTION

a) Constrained Visual-Inertial Localization: [3], and further developments have become popular state-of-the-art methods [8]. These algorithms focus on feature point matching based on binary descriptors, even for small viewpoint changes, where descriptor matching is made real-time applicable by speeding up the process with priorly trained or online learned [9] bag-of-words vocabularies [10, 11].

However, we focus on domains where feature point descriptor matching and tracking perform worse due to challenging image properties [12]. Sparse or semi-dense optical flow [5, 13, 14] are alternatives that perform well forhigh video framerate and can speed up trajectory generation for the *tracking* part of the SLAM system.

Additional information from inertial measurement units (IMU) [15–17] improves visual odometry methods in difficult situations, where tracking is lost, and does not need an external measurement equipment, i.e., avoiding line-of-sight problems.

In contrast to all previously mentioned strategies, we additionally exploit naturally given motion constraints [19], which is also famous in the medical domain as remote center-of-motion (RCM) [20][21], reducing the solution space to tackle the more challenging image domains, which contain large movements in the image, occlusions, fog, deformable surfaces or difficult illuminations conditions [21].

Subsequently, we outline our approach in more detail: We formulate residuals for IMU, stereo camera, and movement constraints and show that the solution space reduces by pivoting the camera around a point from 6 to 4 degrees of freedom. Afterward, we propose to use optical flow in combination with gyroscope readings as an odometry alternative to descriptor-based feature tracking for dynamic image data and high video frequency. Finally, for the refinement of this odometry, we formulate the global cost function and optimization strategies. Our method's main advantage and contribution is the ability to relocalize the camera in unseen environments by minimizing the global cost function, which incorporates all residual terms emerging from sensors

and constraints. The gravity and magnetic field provide the necessary absolute references. In addition, the Lucas-Kanade tracking of the stereo reconstructed locally rigid parts in the environment provides information about the distance from the pivot point.

b) Dataset: Publicly available datasets play a vital role in the evolution of new algorithms, and technologies [22–25]. Researchers focusing on the minimally invasive application have already created endoscopic datasets [21, 25, 27], which are suitable for investigating computer tomography registration, tissue deformation/ reconstruction [28–30], and instrument segmentation. Since they are not complete or are not capturing a handheld intervention, those are unsuitable for testing SLAM applications. In this paper, we introduce a novel dataset suitable for testing SLAM in the MIS domain. To the best of our knowledge, our dataset is the first publicly available surgical dataset comparable to [22–24]. Our MITI Dataset [1], available under the Creative Commons Attribution license (CC-BY 4.0), provides all necessary data by a complete recording of a handheld surgical intervention at Research Hospital Rechts der Isar of TUM. It contains multimodal sensor information from IMU, stereoscopic video, and infrared (IR) tracking as ground truth for evaluation. Furthermore, calibration for the stereoscope, accelerometer, magnetometer, the rigid transformations in the sensor setup, and time-offsets are available. We wisely chose a suitable intervention, namely diagnostic laparoscopy, that contains very few cutting and tissue deformation and shows a full scan of the abdomen with a handheld camera such that it is ideal for testing SLAM algorithms. It incorporates 667 seconds with images (1080p60 RGB) for each camera, readings of two IMU sensors at 220Hz, and IR transformations to two different targets at 20Hz.

Intending to promote the progress of visual-inertial algorithms designed for MIS application, we hope that our clinical training dataset helps and enables researchers to enhance algorithms.

III. NOTATION

The following notions are adapted from [31–33] using the rotation groups $\mathrm{SO}(3)$ and the rigid motion group $\mathrm{SE}(3)$, which fulfill the properties of a Lie group. We use the notation $\exp : \mathfrak{m} \to \mathcal{M}$ to denote the mapping from tangent space, i.e. Lie algebra $\mathfrak{m}$ to Lie group $\mathcal{M}$. The Lie algebra $\mathfrak{so}(3) = \mathrm{skew}(3)$ of $\mathrm{SO}(3)$ has $d=3$ degrees of freedom and $\mathfrak{se}(3)$, tangent space of $\mathrm{SE}(3)$, has $d=6$ degrees of freedom. They are isometric isomorph to $\mathbb{R}^d$, and we use the function Exp as the composition of the hat operator and the matrix exponential $\exp$. The function can explicitly be given in closed form by the Rodriguez formula. The inverse mapping is denoted by Log, which is the mapping from the Lie-Group to $\mathbb{R}^d$. We further use the $\ominus : \mathcal{M} \times \mathcal{M} \to \mathbb{R}^d$ as

$$\mathbf{T}_1 \ominus \mathbf{T}_2 := \mathrm{Log}\left(\mathbf{T}_1^{-1}\mathbf{T}_2\right), \qquad (1)$$

which computes the difference between two transformation matrices mapped to the tangent spaces' isomorph space $\mathbb{R}^d$.

IV. VISUAL-INERTIAL SENSOR FUSION

Our method splits into two parts: Firstly, estimating the pose starting from previous keyframes by movement constraints, gyroscope readings, and Lucas-Kanade tracks of stereoscopic matches. Secondly, the algorithm conducts a windowed optimization over gathered keyframes, changing their location by jointly minimizing the residuals from all modalities, eliminating any drift that might accumulate.

A. Residuals

This subsection describes the residuals gathered from different modalities. The residuals are functions of the localization pose $_\mathrm{T}^\mathrm{C}\mathbf{T} \in \mathrm{SE}(3)$ with the property that they are zero at the actual position and orientation.

a) Motion Constraints: The transformation $_\mathrm{T}^\mathrm{C}\mathbf{T} \in \mathrm{SE}(3)$ transforms a vector in homogeneous coordinates from T to C by multiplication from the right. Since the motion of C is constrained to move around pivot point T but can also move forward and backward in the x-direction, we can reduce the 6 degrees of freedom to 4 degrees of freedom

$$_\mathrm{T}^\mathrm{C}\mathbf{T} = \begin{bmatrix} _\mathrm{T}^\mathrm{C}\mathbf{R} & _\mathrm{T}^\mathrm{C}\boldsymbol{t} \\ \mathbf{0}^T & 1 \end{bmatrix} = \begin{bmatrix} I & _\mathrm{T}^\mathrm{C}\boldsymbol{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} _\mathrm{T}^\mathrm{C}\mathbf{R} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix}, \qquad (2)$$

with $_\mathrm{T}^\mathrm{C}\boldsymbol{t} = t_x \boldsymbol{e}_x$. This means a point in T is first rotated in the pivot point and then translated along the x-axis. The corresponding residual that enforces the result to have only 4 degrees of freedom is

$$\boldsymbol{r}_{pivot} = t_y \boldsymbol{e}_y + t_z \boldsymbol{e}_z. \qquad (3)$$

This adds a penalty to the minimization problem if the solution is not in the reduced space.

b) Accelerometer: The accelerometer readings constrain the current camera pose by one corresponding bearing vector in world coordinates. We assume constant velocity as also proposed by [34]. Rotating the gravity $^\mathrm{T}\boldsymbol{g}$ bearing vector from T to C and comparing it with the measurement $^\mathrm{C}\boldsymbol{g}$ gives the residual

$$\boldsymbol{r}_a = {}^\mathrm{C}\boldsymbol{g} - {}_\mathrm{T}^\mathrm{C}\mathbf{R}\,{}^\mathrm{T}\boldsymbol{g}. \qquad (4)$$

c) Magnetometer: The magnetometer provides similar to the accelerometer a correspondence in world coordinates

$$\boldsymbol{r}_m = {}^\mathrm{C}\boldsymbol{m} - {}_\mathrm{T}^\mathrm{C}\mathbf{R}\,{}^\mathrm{T}\boldsymbol{m}, \qquad (5)$$

with the magnetic field $^\mathrm{T}\boldsymbol{m}$ in T and measurement $^\mathrm{C}\boldsymbol{m}$ in C.

d) Global Mappoints: The distance between a mappoint $^\mathrm{T}\boldsymbol{x}$ in global coordinates and a matched 2D position $\boldsymbol{u}$ in image plane of current camera coordinates is

$$\boldsymbol{r}_{2D} = \Pi\big(_\mathrm{T}^\mathrm{C}\mathbf{T}\,{}^\mathrm{T}\boldsymbol{x}\big) - \boldsymbol{u} \qquad (6)$$

with $\Pi(\cdot)$ being the projection into distorted image plane and $\boldsymbol{u} := \boldsymbol{u}_d$ the distorted image coordinates. In (6), the map points $^\mathrm{T}\boldsymbol{x}$ are connected rigidly to the world coordinate system T at initialization by using the current camera pose estimate. We need to rewrite (6), attaching the mappoints to the camera coordinate system. Otherwise, since we optimize only the poses $_\mathrm{T}^\mathrm{C}\mathbf{T}$, the minimization of other residuals leads to the increase of (6). For this sake, we store the landmark

positions relative to the camera coordinate system $\texttt{C}$, in which the stereo matcher initialized them. This adds another computation step and one more parameter ${}_{\texttt{C}}^{\texttt{T}}\mathbf{T}_{t'}$ at detection time t' to the residual in (6), namely

$$ {}^{\texttt{T}}\boldsymbol{x} = {}_{\texttt{C}}^{\texttt{T}}\mathbf{T}_{t'}\,{}^{\texttt{C}}\boldsymbol{x}_{t'}. \tag{7} $$

Thus we have

$$ \boldsymbol{r}_{2D} = \Pi({}_{\texttt{T}}^{\texttt{C}}\mathbf{T}\,{}_{\texttt{C}}^{\texttt{T}}\mathbf{T}_{t'}\,{}^{\texttt{C}}\boldsymbol{x}_{t'}) - \boldsymbol{u}, \tag{8} $$

which connects the observation frame with the initialization frame of a landmark.

e) Gyroscope: The gyroscope measurements in tangent space of $\boldsymbol{\omega}_{t'} \in \mathbb{R}^3 \equiv \mathfrak{so}(3)$ are integrated by gyroscope sample time Δt and the non-abelsh product

$$ {}^{\texttt{C}}\Delta\mathbf{R}_{t_1,t_2} := \prod_{t' \in \{t_1,\dots,t_2-\Delta t\}} \mathrm{Exp}({}^{\texttt{C}}\boldsymbol{\omega}_{t'}\Delta t). \tag{9} $$

This gives the update rule

$$ {}_{\texttt{C}}^{\texttt{T}}\mathbf{R}_{t_2} = {}_{\texttt{C}}^{\texttt{T}}\mathbf{R}_{t_1}\,{}^{\texttt{C}}\Delta\mathbf{R}_{t_1,t_2} \tag{10} $$

and the residual term

$$ \boldsymbol{r}_g = {}_{\texttt{C}}^{\texttt{T}}\mathbf{R}_{t_2} \ominus {}_{\texttt{C}}^{\texttt{T}}\mathbf{R}_{t_1}\,{}^{\texttt{C}}\Delta\mathbf{R}_{t_1,t_2}. \tag{11} $$

which constraints consecutive transformations to have a certain deviation in orientation, namely the integrated measurements of the angular velocity.

B. Visual-Inertial Odometry Tracking

This subsection describes our approach for the odometry tracking thread. It updates from the pose of the previous keyframe, generating good initializations for the optimization using gyroscope data and Lucas Kanade tracks of stereo correspondences. Algorithm 1 sketches the procedure, and in the following, we will explain line by line the computation steps the algorithm conducts.

After entering the loop, it first updates the pose based on the gyroscope measurements in Line 2. Although (11) usually only provides information about the orientation, a position update is possible due to the camera movement having only 4 degrees of freedom

$$ {}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_2} = \begin{bmatrix} I & {}_{\texttt{T}}^{\texttt{C}}\boldsymbol{t}_{t_1} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} {}^{\texttt{C}}\Delta\mathbf{R}_{t_1,t_2}^T\,{}_{\texttt{T}}^{\texttt{C}}\mathbf{R}_{t_1} & \mathbf{0} \\ \mathbf{0}^T & 1 \end{bmatrix} \tag{12} $$

with ${}_{\texttt{T}}^{\texttt{C}}\boldsymbol{t}_t = t_{x,t}\boldsymbol{e}_x$. The new pose is on a sphere with radius $t_{x,t}$, around the trocar entry point, leaving only the radius unknown but estimating all other degrees of freedom.

During the first iteration, the Lines 3 to 4 have no effect. The first frame is a keyframe, and the algorithm computes stereo correspondences by descriptor matching using ORB descriptors [35] in Line 6. ORB descriptors are faster to compute and compare than SIFT or SURF, whose benefits, e.g., scale and rotation invariance, are less essential for stereo-matching. For faster matching, it uses the projections provided by the camera calibration parameters as initialization for finding matches in the second image [8]. The epipolar error threshold [36] is adjusted to filter outliers

$$ {}^{\texttt{C0}}\boldsymbol{x}^T \mathbf{E}\,{}^{\texttt{C1}}\boldsymbol{x} > d_{th}, \tag{13} $$

with epipolar matrix $\mathbf{E} = \hat{\boldsymbol{t}}\mathbf{R}$ and $\hat{\cdot}: \mathbb{R}^3 \to \mathrm{skew}(3)$, defining the transformation between left $\texttt{C0}$ and right $\texttt{C1}$ lens. The matched points initialize new landmarks in Line 7.

Afterwards, in Line 3 for every incoming frame it updates the landmark projections $\boldsymbol{u} \in \mathbb{R}^2$ in the 2D image plane, producing residuals of the form (6). For that, it uses the variational Lucas-Kanade-Method [37] taking two consecutive frames for every camera. This variational method performs well (small baseline, photometric consistency) if video frequency is high. For large displacements, the drift becomes large, and for fast movements, we lose track. Nevertheless, optical flow is suitable for local tracking.

With initialization computed in Line 2 it determines the camera pose ${}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_t$ based on the 3D-2D-correspondences in Line 4 by a perspective-n-point algorithm [38] inside a random sample consensus (RANSAC) [39] scheme with refinement provided by OpenGV [40]. A sanity check on the rotation part using the IMU measurements and on the translation part using the constraints ensures locally rigid features.

Algorithm 1 Visual-Inertial Odometry with Optimization

Input: Initialization ${}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_0}$, Images $\{I_{0,t_i}, I_{1,t_i}\}_{i \in I}$;
　　IMU Data $\{{}^{\texttt{C}}\Delta\mathbf{R}_{t_i,t_{i+1}}, {}^{\texttt{C}}\boldsymbol{a}_{t_i}, {}^{\texttt{C}}\boldsymbol{m}_{t_i}\}_{i \in I}$
Output: $\{{}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_i}\}_{i \in I}$
1: **while** True **do**
2:　　Update ${}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_{i-1}}$ based on ${}^{\texttt{C}}\Delta\mathbf{R}_{t_{i-1}}$ and movement constraints
3:　　Lukas-Kanade tracking of landmarks
4:　　Update ${}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_i}$ based on 2D-landmark projections
5:　　**if** Decision for keyframe **then**
6:　　　Stereo matching
7:　　　Landmark insertion
8:　　**end if**
9:　　$i = i + 1$
10: **end while**

C. Optimization

The optimization presented in this subsection jointly minimizes all previously formulated residuals over a window of keyframes. As shown later the camera can globally relocalize itself by solving the proposed optimization problem solely by small tracks of locally rigid feature points and IMU measurements. In order to use all the information from sensor readings and images we build a cost function by summing up the normalized, weighted and robustified residuals that have been collected, i.e. (3, 4, 5, 8, 11) which are available at each timestep. We denote them as residuals $\boldsymbol{r}_k$, $k \in K$, where k indicates time and residual type.

The overall optimization problem is

$$ \underset{{}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{n:m}}{\arg\min} \sum_k \rho_k\left(\|\alpha_k \boldsymbol{r}_k\left({}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{n:m}\right)\|\right) := \sum_k f_k({}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{n:m}), \tag{14} $$

optimizing over the window of keyframe localizations ${}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{n:m} := ({}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_n}, \dots, {}_{\texttt{T}}^{\texttt{C}}\mathbf{T}_{t_m})$. In above formulation, α_k stands

for the scaling factor, and ρ for the robustification function. Each residual gets an individual scaling factor α_k. The evaluation part of this paper gives a more detailed overview of the chosen weight factors, optimization window size, and trigger of the optimization problem.

For robustification of residuals that are prone to outliers, we choose the Huber loss

$$\rho_k(x) = \begin{cases} \frac{1}{2}x^2 & \text{for } |x| \le \delta, \\ \delta(|x| - \frac{1}{2}\delta), & \text{otherwise.} \end{cases} \quad (15)$$

The function input x in our application is the normalized and weighted residual $\|\alpha_k r_k\|$. For residuals with a dense error distribution and no outliers, ρ_k is simply the identity.

Each of the residuals $r_k\left(\,^c_{\mathrm{T}}\mathbf{T}_{n:m}\right)$ depends only on a maximum of two parameters $^c_{\mathrm{T}}\mathbf{T} \in \mathrm{SE}(3)$, and further, many parameters do not have a residual in common. The residual terms give the optimization problem the typical graph structure of SLAM problems. This sparse dependence on parameters makes the Hesse-matrix sparse, and the optimization problem suitable for a sparse Schur solver [41].

V. DATASET

In this section, we introduce our dataset and calibration procedures. In Fig. 1 we show the sensor setup during the data acquisition process. The sensor data consists firstly of a stereoscopic video (Tipcam1, Karl Storz) from the view inside the abdomen. Secondly, two 9DOF IMU (MetamotionR, Metawear) sensor boards attached to the laparoscope handle sending: angular velocity, acceleration, and magnetic field strength. Thirdly, an IR tracking camera (Polaris Vega, NDI) detects two targets, each consisting of 3-4 reflecting passive spheres. The anticipated precision of the IR tracking is 0.25mm [42], making it suitable as ground truth to evaluate the SLAM algorithm's performance. We calibrated the camera and the IMU, determining sensor-specific parameters and calculated time offsets using the state-of-the-art methods defined in subsequent subsections. Parameters can be found in the path `calibration/*.csv` of the dataset [1].

A. Camera Calibration

The pinhole projection [43] $\Pi_0 : \mathbb{R}^3 \to \mathbb{R}^2$ $x \mapsto u_u$ from a point $x \in \mathbb{R}^3$ in coordinate system of left `C0` or right `C1` lens to the undistorted image point u_u in the image plane is parameterized by focal lengths f_x, f_y and optical center c_x, c_y both measured in pixel.

Radial image distortion is compensated by a low order polynomial model $d : \mathbb{R}^2 \to \mathbb{R}^2$, $u_u \mapsto u_d$. Additionally,
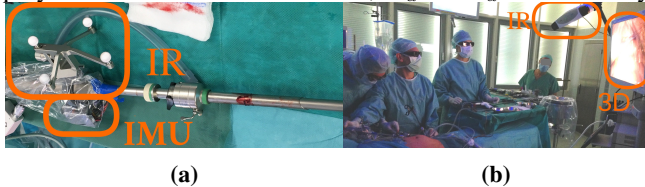

(a)   (b)

Fig. 1: Data Acquisition (a) IMU sensorboards and IR passive targets attached to stereoscopic laparocopic camera (b) IR sensor and 3D monitor in the operating room
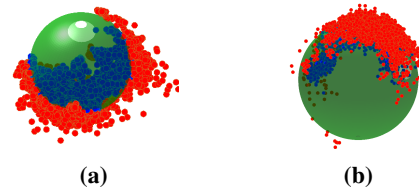

(a)   (b)

Fig. 2: Calibrated IMU measurements, i.e., red in front, blue behind sphere (a) $^c a - {}^c g$ in $9.81 \frac{m}{s^2}$ and unit sphere (b) $^c m$ in μT and sphere with radius 48.6

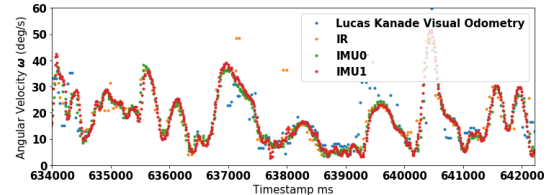**Fig. 3: Time-calibrated angular velocity part $\|\,^c\omega_{t_i}^{(k)}\|$**

the extrinsic parameters are calibrated, which determine the rigid body transformation $^{\mathrm{C0}}_{\mathrm{C1}}T \in \mathrm{SE}(3)$. The function $\Pi(\cdot)$ in residual (6) is a composition of rotation and translation from `C` to `C0` or `C1` and projection to distorted 2D image coordinates by composition of Π_0 and d. (**??**).

B. IMU Calibration

Using a standard approach [34, 44], the accelerometer measurements $\tilde{a}(t)$ are modeled as

$$\tilde{a}(t) = (a(t) - g(t))d_a^{-1} + n_a(t) + b_a \quad (16)$$

with acceleration $a(t) \in \mathbb{R}^3$, gravity $g(t) \in \mathbb{R}^3$, errors $b_a \in \mathbb{R}^3$, $d_a \in \mathbb{R}$ and $n_a(t) \sim \mathcal{N}(0, \sigma_a^2 \mathbf{I})$. Calibration assumes constant velocity and gravity norm $g = \|g(t)\| = 9.81 \frac{m}{s^2}$. Furthermore, we model the error of magnetometer readings $\tilde{m}(t)$ [34, 44] by

$$\tilde{m}(t) = m(t)d_m^{-1} + n_m(t) + b_m, \quad (17)$$

with magnetic field strength $m(t) \in \mathbb{R}^3$, errors $b_m \in \mathbb{R}^3$, $d_m \in \mathbb{R}$ and $n_m(t) \sim \mathcal{N}(0, \sigma_m^2 \mathbf{I})$. The corrected measurements lie on a sphere with radius equal to the strength of the earths magnetic field $m = \|m\| = 48.6 \mu T$ in Munich, Germany. This information is used in the calibration procedure. Fig. 2 shows calibrated sensor readings.

C. Time Offset Calibration

The sensor data has been synchronized during the acquisition process by configuring a network time protocol server/client at the data receiving devices. For synchronization of the video the displayed timestamp has been captured at the start and end of the intervention. The sensor interfaces for IR/IMU are ethernet/bluetooth respectively, which causes a remaining timeoffset in the collected data. This offset is minimized by evaluating the tangent $^c\xi_{t_i} \in \mathbb{R}^6 \cong \mathfrak{se}(3)$ for each sensor at coordinate system `C` at time t_i We then calibrated the timeoffset dt by comparing the velocities acquired from different sensors

$$j, k \in \{\text{"Visual Odometry", "IR", "IMU0", "IMU1"}\}.$$
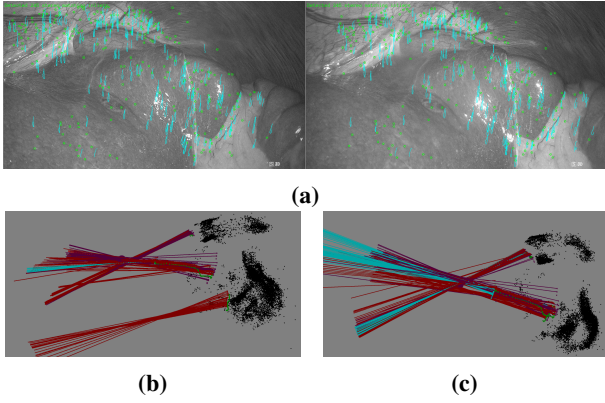
(a)

(b) **(c)**

Fig. 4: Debug environment (a) Stereo view with ORB-descriptor-based stereomatches in green, Lucas-Kanade-tracks in teal, (b) and (c) Camera poses connected to lines, that should go through pivot point and 3D map points of left and right diaphragm, (b) Camera-only based tracking, (c) Incorporating IMU sensordata and movement constraint

The sum of differences at different timeoffsets is

$$\min_{dt} \sum_i \left\| {}^{c}\boldsymbol{\xi}_{t_i+dt}^{(j)} - {}^{c}\boldsymbol{\xi}_{t_i}^{(k)} \right\|. \tag{18}$$

For that, we determine the camera-based angular velocity using the visual odometry algorithm based on Lucas Kanade feature tracks without IMU and IR. The computed time-offsets are 100ms/166ms for the IR/IMU readings. Time synchronization is vital for short-term odometry, unlike windowed optimization, where keyframes are more distant.

VI. EVALUATION AND EXPERIMENTS

We conduct experiments on an implementation of the approach described in Section IV. In Subsection VI-C we examine the influence of weight distribution on residuals, and in Subsection VI-B we discuss the advantages of utilizing robust loss functions. Subsection VI-A and VI-D are devoted to the evaluation of our method with our MITI dataset. In doing that, we measure the deviation of the algorithms output to ground-truth poses obtained by IR tracking.

We used several libraries: OpenCV for keypoint detection, feature tracking, and reprojection; OpenGV [40] for triangulation, RANSAC, and nonlinear refinement; as well as Ceres [45] and Sophus [46] for solving the optimization problem. For the experiment, we took the image sequence starting from timestamp 637273ms, which contains a fast movement from left to right with dynamic occlusions and enables us to demonstrate the algorithm's performance in a scenario where standard SLAM approaches lose tracking and can not relocalize in the unseen environment.

A. Evaluation of Visual Odometry Tracking

In Fig. 4 one can inspect the debug view of the output and Fig. 4b and Fig. 4c show huge deviations for different settings in the algorithm. The isolated bundle of lines in Fig. 4b shows: without movement constraints and IMU measurements, the optimization is unable to converge to
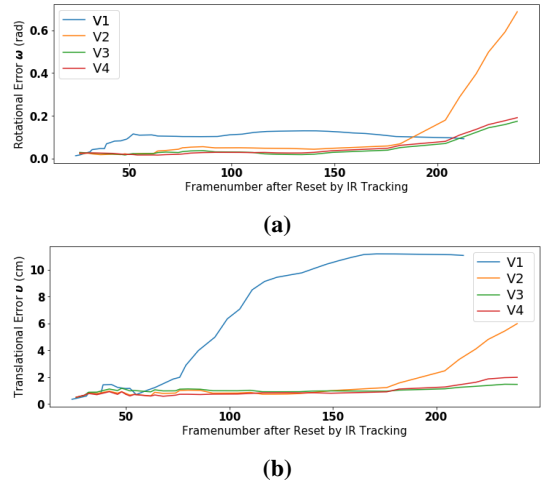

(a)

(b)

Fig. 5: Algorithm 1 without optimization (a) V1 relies on IMU data and can reduce rotational error in the long run, whereas V2 based on vision suffers from drift, (b) error of V1 is unacceptable when movements do not stay on the sphere, error of pure visual odometry V2 also suffers from drift

the correct poses since it loses tracking of 2D features in between keyframes. We computed the Log distance between IR tracked and constrained visual-inertial odometry (VO) estimated camera poses

$$\boldsymbol{\xi} = [\boldsymbol{\omega}, \boldsymbol{v}]^T = {}^{c}_{T}\mathbf{T}^{(IR)} \ominus {}^{c}_{T}\mathbf{T}^{(VO)} \in \mathbb{R}^6 \tag{19}$$

for variants of the computation steps in Algorithm 1, that are listed in Table I. V1 integrates IMU sensor-readings in Line 2 into pose estimates as given in (12). Therefore, it cannot detect changes in the x-translation component, resulting in a large translational error as soon as the camera moves inside. V2 only uses the camera-based pose estimation of Lines 3 and 4 but without taking the initialization provided by the IMU in Line 2. The result of this variant can additionally be inspected in the debug view of Fig. 4b. V3 uses both camera-based and IMU-based tracking, which already leads to a significant decrease in both error residuals. Finally,

TABLE I: Settings of Alg. 1 for results shown in Fig. 5

Variant	Pivot point online	Line2	Line 3 and 4	Optimization
V1	×	✓	×	×
V2	×	×	✓	×
V3	×	✓	✓	×
V4	✓	✓	✓	×

in V4, instead of considering the pivot point given by the priorly computed position from IR tracking, we initialize and iteratively refine the pose purely based on visual odometry. The result is an essential step for the goal an eternal inside-out tracking.

B. Error Analysis and Choice of Robustification Functions

The residuals themselves are composed of errors in sensor readings and the deviation of the current pose estimate. Aiming to minimize only the part indicating a pose deviation, we examine errors of sensor readings and their propagation

TABLE II: Statistic of Residuals

	Eq.	Expectation $\mathrm{E}[\boldsymbol{r}]$	Variance $\mathrm{Var}[\boldsymbol{r}]$	Unit
$\boldsymbol{r}_{2D}$	(6)	1.6263	9.3767	Pixel
$\boldsymbol{r}_g$	(9)	0.0003	0.0051	rad
$\boldsymbol{r}_{pivot}$	(3)	0.0511	0.4374	cm
$\boldsymbol{r}_a$	(4)	0.0235	0.0405	$9.81\frac{m}{s^2}$
$\boldsymbol{r}_m$	(5)	-0.2730	4.2801	μT

to the residual term to decide whether to use a Huber loss. The Huber loss for the reprojection residual (8) is justified, since we identified an outlier characteristic in the three error sources: A drift error in the Lucas-Kanade tracked 2D locations $^c\boldsymbol{u}$, that can be modeled as a random walk, after n tracking steps, contributes with a Gaussian probability distribution of $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, n\sigma^2\mathbf{I})$ linearly to the residual. The second error source is a badly reconstructed landmark position $^c\boldsymbol{x}_{t'}$, which occurs although the epipolar error threshold reduces the probability of outliers. The third error arises from a failing filtering of dynamic feature points. For residuals that penalize solutions not going through the trocar entry point (3) the use of the Huber loss is not justified, since there is no measurement error, and we aim to minimize large residuals. The error of IMU measurements (4, 5, 9) is approximately Gaussian, which leads us to not choosing the Huber loss.

C. Weighting Residuals

In this subsection we discuss the factor α_k that strongly influences the optimization (14). Firstly, a task of the parameter is to adjust the residuals, making them independent on their measurement unit. We ran Algorithm 1 and built a statistic for every residual type, as shown in Tab II.

The expectation of the residuals is close to zero, which simplifies the normalization $\frac{\boldsymbol{r}-\mathrm{E}[\boldsymbol{r}_k]}{\sqrt{\mathrm{Var}[\boldsymbol{r}_k]}}$. The scaling factor splits into $\alpha_k = \frac{1}{N_k}\frac{1}{\sqrt{\mathrm{Var}[\boldsymbol{r}_k]}}\beta_k$, with β_k being an additional term to adjust the importance between residuals, as discussed in the next Subsection. Hereby, we additionally normalize by the number of residuals of the same type N_k e.g., there exist $N_k = 1$ gyroscope readings per keyframe, but in general $N_k > 1$ reprojection residuals per keyframe.

D. Evaluation of Trade-off and Optimization

A trade-off between gyroscope and camera-based residuals is adjusted by $\beta_k = \gamma$ for residual k belonging to gyroscope measurements and $\beta_k = (1 - \gamma)$ for reprojection residuals, with $\gamma \in [0, 1]$. We conduct experimental runs with optimization being triggered every 10 keyframes and a window size of 10. Fig. 6a and 6b show the deviation of the SLAM algorithm from the ground-truth IR tracked poses for different values of γ. Since we want to jointly optimize over all residuals, the cost emerging from different sensors should simultaneously decrease. The cost function difference Δf is

$$\Delta f = f(^c_T\mathbf{T}_{n:m}^{(\text{after})}) - f(^c_T\mathbf{T}_{n:m}^{(\text{before})}), \qquad (20)$$

with evaluations before and after optimization. Negative values indicate that this k-th part of the cost function was minimized. In Fig. 6c we show the average cost difference for different residual types in one optimization step for different

(a)

(b)

(c)
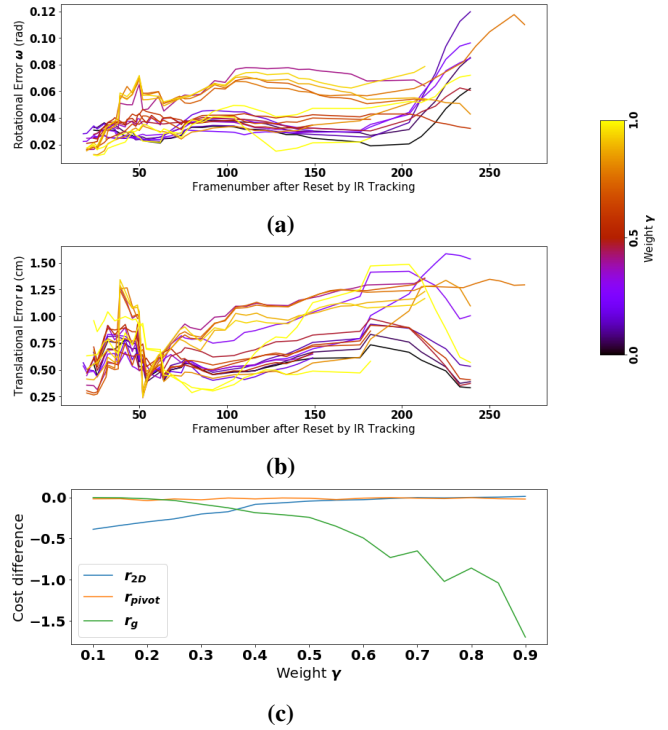
Fig. 6: Errors and cost reduction of V3 with optimization for different values of tradeoff value γ

values of γ. It exhibits a simultaneous cost decrease for values close to $\gamma = 0, 4$.

On the other hand, we observe by inspecting the error trajectories in Fig. 6a and 6b that for values close to 0.5 the optimization successfully incorporates the advantages of both sensors: The rotational error is low at the beginning, which is typical for camera-based tracking, but also stays down, which is possible due to utilizing the information of IMU sensors. The distribution of weights shows great potential in influencing the overall outcome of the algorithm. An optimal adjustment however, is a dynamic weight distribution, depending on the current conditions that influence the reliability of the sensor data.

VII. DISCUSSION

Movement constraints and additional sensor information are necessary to provide stable pose estimates in dynamic environments. Absolute references are provided by the earth's gravity and magnetic field, whereas gyroscope and Lucas-Kanade tracks provide the odometry. For a more robust setup, however, descriptor matching of rigid features can improve the results of long-term SLAM.

VIII. CONCLUSIONS

This work shows that visual-inertial algorithms can be applied to dynamically challenging fields such as minimally invasive surgery. This sets a new paradigm in the applicability of SLAM. The gain of reliability is due to the fact that we use inherent motion constraints. Our results show for the use-case of minimally invasive surgery, that self localization is indeed feasible and hope that researchers can profit from our publicly available dataset.

References

[1] R. Hartwig, D. Ostler, J.-C. Rosenthal, H. Feussner, D. Wilhelm, and D. Wollherr, "Miti: Slam benchmark for laparoscopic surgery," TUM, 2021. [Online]. Available: https://mediatum.ub.tum.de/1621941

[2] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *Trans. on robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.

[3] G. Klein and D. Murray, "Parallel tracking and mapping for small ar workspaces," in *Int. symposium on mixed and augmented reality*, 2007, pp. 225–234.

[4] J. Janai, F. Güney, A. Behl, A. Geiger *et al.*, "Computer vision for autonomous vehicles: Problems, datasets and state of the art," *Foundations and Trends in Computer Graphics and Vision*, vol. 12, no. 1–3, pp. 1–308, 2020.

[5] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "Svo: Semidirect visual odometry for monocular and multicamera systems," *Trans. on Robotics*, vol. 33, no. 2, pp. 249–265, 2016.

[6] C. Kerl, J. Sturm, and D. Cremers, "Robust odometry estimation for rgb-d cameras," in *Int. Conf. on Robotics and Automation (ICRA)*, 2013.

[7] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, "Keyframe-based visual–inertial odometry using nonlinear optimization," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 314–334, 2015.

[8] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *Trans. on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.

[9] S. Khan and D. Wollherr, "Ibuild: Incremental bag of binary words for appearance based loop closure detection," in *Int. Conf. on Robotics and Automation (ICRA)*, 2015, pp. 5441–5447.

[10] D. Gálvez-López and J. D. Tardos, "Bags of binary words for fast place recognition in image sequences," *Trans. on Robotics*, vol. 28, no. 5, pp. 1188–1197, 2012.

[11] D. Nister and H. Stewenius, "Scalable recognition with a vocabulary tree," in *Int. Conf. on Computer Vision and Pattern Recognition*, vol. 2. IEEE, 2006, pp. 2161–2168.

[12] L. Maier-Hein, A. Groch, A. Bartoli, S. Bodenstedt, G. Boissonnat, P.-L. Chang, N. T. Clancy, D. S. Elson, S. Haase, E. Heim, J. Hornegger, P. Jannin, H. Kenngott, T. Kilgus, B. Müller-Stich, D. Oladokun, S. Röhl, T. R. dos Santos, H.-P. Schlemmer, A. Seitel, S. Speidel, M. Wagner, and D. Stoyanov, "Comparative validation of single-shot optical techniques for laparoscopic 3-d surface reconstruction," *Trans. on Medical Imaging*, vol. 33, no. 10, pp. 1913–1930, 2014.

[13] Q. Fu, H. Yu, X. Wang, Z. Yang, H. Zhang, and A. Mian, "Fastorb-slam: A fast orb-slam method with coarse-to-fine descriptor independent keypoint matching," *arXiv preprint arXiv:2008.09870*, 2020.

[14] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *Trans. on pattern analysis and machine intelligence*, vol. 40, no. 3, pp. 611–625, 2017.

[15] L. Von Stumberg, V. Usenko, and D. Cremers, "Direct sparse visual-inertial odometry using dynamic marginalization," in *Int. Conf. on Robotics and Automation (ICRA)*, 2018, pp. 2510–2517.

[16] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *Trans. on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.

[17] T. Qin, J. Pan, S. Cao, and S. Shen, "A general optimization-based framework for local odometry estimation with multiple sensors," *arXiv preprint arXiv:1901.03638*, 2019.

[18] C. Mei, G. Sibley, M. Cummins, P. Newman, and I. Reid, "Rslam: A system for large-scale mapping in constant-time using stereo," *Int. journal of computer vision*, vol. 94, no. 2, pp. 198–214, 2011.

[19] M. Salerno, K. Zhang, A. Menciassi, and J. S. Dai, "A novel 4-dofs origami enabled, sma actuated, robotic end-effector for minimally invasive surgery," in *Int. Conf. on Robotics and Automation (ICRA)*, 2014, pp. 2844–2849.

[20] F. Vasconcelos, E. Mazomenos, J. Kelly, and D. Stoyanov, "Rcm-slam: Visual localisation and mapping under remote centre of motion constraints," in *2019 International conference on robotics and automation (ICRA)*. IEEE, 2019, pp. 9278–9284.

[21] B. Lin, Y. Sun, X. Qian, D. Goldgof, R. Gitlin, and Y. You, "Video-based 3d reconstruction, laparoscope localization and deformation recovery for abdominal minimally invasive surgery: a survey," *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 12, no. 2, pp. 158–178, 2016.

[22] S. Baker, D. Scharstein, J. Lewis, S. Roth, M. J. Black, and R. Szeliski, "A database and evaluation methodology for optical flow," *Int. journal of computer vision*, vol. 92, no. 1, pp. 1–31, 2011, available at https://vision.middlebury.edu/flow/. [Online]. Available: https://vision.middlebury.edu/flow/

[23] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Int. Conf. on computer vision and pattern recognition*, 2012, pp. 3354–3361.

[24] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *Int. Conf. on Intelligent Robots and Systems*, 2012, pp. 573–580.

[25] M. Allan, J. Mcleod, C. Wang, J. C. Rosenthal, Z. Hu, N. Gard, P. Eisert, K. X. Fu, T. Zeffiro, W. Xia *et al.*, "Stereo correspondence and reconstruction of endoscopic data challenge," *arXiv preprint arXiv:2101.01133*, 2021.

[26] M. D. Blikkendaal, S. R. Driessen, S. P. Rodrigues, J. P. Rhemrev, M. J. Smeets, J. Dankelman, J. J. van Den Dobbelsteen, and F. W. Jansen, "Surgical flow disturbances in dedicated minimally invasive surgery suites: an observational study to assess its supposed superiority over conventional suites," *Surgical endoscopy*, vol. 31, no. 1, pp. 288–298, 2017.

[27] D. Dey, D. Gobbi, P. Slomka, K. Surry, and T. Peters, "Automatic fusion of freehand endoscopic brain images to three-dimensional surfaces: creating stereoscopic panoramas," *Trans. on Medical Imaging*, vol. 21, no. 1, pp. 23–30, 2002.

[28] N. Mahmoud, T. Collins, A. Hostettler, L. Soler, C. Doignon, and J. M. M. Montiel, "Live tracking and dense reconstruction for handheld monocular endoscopy," *Trans. on Medical Imaging*, vol. 38, no. 1, pp. 79–89, 2019.

[29] P. Mountney, D. Stoyanov, and G.-Z. Yang, "Three-dimensional tissue deformation recovery and tracking," *Signal Processing Magazine*, vol. 27, no. 4, pp. 14–24, 2010.

[30] L. Maier-Hein, P. Mountney, A. Bartoli, H. Elhawary, D. Elson, A. Groch, A. Kolb, M. Rodrigues, J. Sorger, S. Speidel *et al.*, "Optical techniques for 3d surface reconstruction in computer-assisted laparoscopic surgery," *Medical image analysis*, vol. 17, no. 8, pp. 974–996, 2013.

[31] T. D. Barfoot, *State estimation for robotics*. Cambridge University Press, 2017.

[32] J. Sola, J. Deray, and D. Atchuthan, "A micro lie theory for state estimation in robotics," *arXiv preprint arXiv:1812.01537*, vol. 5, p. 1371, 2018.

[33] C. Sommer, V. Usenko, D. Schubert, N. Demmel, and D. Cremers, "Efficient derivative computation for cumulative b-splines on lie groups," in *Int. Conf. on Computer Vision and Pattern Recognition*, 2020, pp. 11 148–11 156.

[34] S. Madgwick, "An efficient orientation filter for inertial and inertial/magnetic sensor arrays," *Report x-io and University of Bristol (UK)*, vol. 25, pp. 113–118, 2010.

[35] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *Int. Conf. on computer vision*. IEEE, 2011, pp. 2564–2571.

[36] D. Forsyth and J. Ponce, *Computer vision: A modern approach*. Prentice hall, 2011.

[37] A. Bruhn, J. Weickert, and C. Schnörr, "Lucas/kanade meets horn/schunck: Combining local and global optic flow methods," *Int. journal of computer vision*, vol. 61, no. 3, pp. 211–231, 2005.

[38] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Epnp: An accurate o (n) solution to the pnp problem," *International journal of computer vision*, vol. 81, no. 2, pp. 155–166, 2009.

[39] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[40] L. Kneip and P. Furgale, "Opengv: A unified and generalized approach to real-time calibrated geometric vision," in *Int. Conf. on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 1–8.

[41] G. Grisetti, R. Kümmerle, H. Strasdat, and K. Konolige, "g2o: A general framework for (hyper) graph optimization," in *Int. Conf. on Robotics and Automation (ICRA)*, 2011, pp. 9–13. [Online]. Available: http://ais.informatik.uni-freiburg.de/publications/papers/kuemmerle11icra.pdf

[42] G. Fattori, A. J. Lomax, D. C. Weber, and S. Safai, "Technical assessment of the ndi polaris vega optical tracking system," *Radiation oncology*, vol. 16, no. 1, pp. 1–4, 2021.

[43] Z. Zhang, "A flexible new technique for camera calibration," *Trans. on pattern analysis and machine intelligence*, vol. 22, no. 11, pp. 1330–1334, 2000.

[44] V. Renaudin, M. H. Afzal, and G. Lachapelle, "Complete triaxis magnetometer calibration in the magnetic domain," *Journal of sensors*, vol. 2010, pp. 1–10, 2010.

[45] S. Agarwal and K. Mierle, "Ceres solver: Tutorial & reference," *Google Inc*, vol. 2, no. 72, p. 8, 2012.

[46] H. Strasdat, S. Lovegrove, and Others, "Sophus library," https://github.com/strasdat/Sophus.