

Joint Axis and Position Estimation from Inertial Measurement Data by Exploiting Kinematic Constraints

Thomas Seel, Thomas Schauer, Jörg Raisch

Abstract—We consider 6d inertial measurement units (IMU) attached to rigid bodies, e.g. human limb segments or links of a robotic manipulator, which are connected by hinge joints and spheroidal joints. Novel methods for joint axis estimation and joint position estimation are presented that exploit the kinematic constraints induced by these two types of joints. The presented methods do not require any knowledge about the sensor units' positions or orientations and do not include integration, i.e. they are insensitive to measurement bias. By means of a three-links simulation model, the estimation algorithms are validated and convergence is analyzed. Finally, the algorithms are tested using experimental data from IMU-based human gait analysis.

I. INTRODUCTION

Inertial measurement units (IMU) have been used in a multitude of applications and have recently caught increasing attention in the field of human motion analysis. Typically, each limb segment is equipped with one IMU and the goal is to measure position and orientation of the links as well as one, two or three joint angles, in case of a hinge, saddle or spheroidal joint, respectively. Apparently, this information can be gained from the measured accelerations and angular velocities by strap-down integration and some coordinate transformation. However, two fundamental problems arise: One is that since even highly accurate sensors do not have a bias of exactly zero, position and angle estimates are always subject to drift. And the second is that information is required about the (constant) orientation of the sensor's coordinate systems with respect to the joint axis or the segments they are mounted on.

Regarding the drift problem, many suggestions have been made on how this effect might be overcome, but most of them are based on additional assumptions like periodic phases of rest [7], an even ground [6], or a constant magnetic field [2]. Besides these rather restrictive approaches there is also considerable work on reducing drift effects by exploiting the kinematic constraints that the segments of human limbs or a robotic manipulator are bound to, see e.g. [1]. And in [5], kinematic constraints are used to calculate joint angles without requiring any integration. In both cases, however, the offset vectors from the origins of the sensor frames to the joint center are assumed to be known. And even when using more than one IMU per segment, as in

[8], exactly this information is crucial.

Regarding the orientation problem, a number of previous approaches, e.g. [3] and [2], used calibration movements in order to estimate the orientation of the local sensor frame with respect to a global frame or the joint axis. Apparently, these techniques can lead to very poor results unless a tight mechanical setup is used to restrict the motion. A tempting alternative is to mount the sensors with a predefined orientation towards the segment or joint, as in [4] and [5]. But besides the fact that this is hard to realize for some applications, e.g. human gait analysis, the casing of a sensor unit rarely coincides with its inner coordinate system.

Therefore, there is high demand for methods that enable accurate estimation of joint positions and joint axes with respect to the local sensor frames. In this contribution we will demonstrate how this information can be extracted from the measurement data of almost arbitrary movements by exploiting the kinematic constraints of the respective joints. Section II explains how the reduction in degrees of freedom become manifest in the accelerations and angular velocities of the sensor units. A geometric model is introduced in Section III in order to provide a proof of concept and more insight into the underlying ideas. Finally, in Section IV, we present experimental results based on inertial measurement data from a human gait analysis.

II. EXPLOITING KINEMATIC CONSTRAINTS

A. Constraints induced by hinge joints

Consider two rigid segments that are free to rotate and move in space but are connected by a hinge joint, as depicted in Fig. 1. The segments shall be called the first and the second segment, and each of them shall be equipped with a three-dimensional gyroscope that is attached to the segment in some arbitrary orientation. The unit joint axis vector with respect to the local coordinate system of the first segment's gyroscope shall be referred to as j_1 . And j_2 shall be the very same unit joint axis vector but seen from the local coordinate system of the second segment's gyroscope. Moreover, let the angular velocities of the gyroscopes, in the coordinates of their local frames, be $g_1(t)$ and $g_2(t)$ for the first and the second segment, respectively. Then it is a geometrical fact, that $g_1(t)$ and $g_2(t)$ differ only by the joint angle velocity and a (time-variant) rotation matrix. Hence their projections into the joint plane¹ have the same lengths for each instant

¹i.e. the plane to which the joint axis is the normal vector

T. Seel and T. Schauer are with Control Systems Group at TU Berlin, Germany, seel@control.tu-berlin.de

Jörg Raisch is with Control Systems Group at TU Berlin and Systems and Control Theory Group at Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany

in time, which is equivalent to

$$\|g_1(t) \times j_1\|_2 - \|g_2(t) \times j_2\|_2 = 0 \quad \forall t, \quad (1)$$

where $\|\cdot\|_2$ denotes the Euclidean norm. This fact turns out to be very useful when faced with the task of identifying the hinge joint axis in case the orientation of the sensors towards the segments is unknown. One can simply choose a large set of gyroscopic data from both sensors and search for the vectors \hat{j}_1 and \hat{j}_2 that fulfill (1) in a least-squares sense. Therefore, we calculate the gradients of the left-hand side terms of (1) with respect to j_1 and j_2 :

$$\frac{d(\|g_i(t) \times j_i\|_2)}{dj_i} = \frac{(g_i(t) \times j_i) \times g_i(t)}{\|g_i(t) \times j_i\|_2}, \quad i = 1, 2. \quad (2)$$

This can be used, e.g., in a Gauss-Newton algorithm, as further explained in Section III-B. Of course, the motion of the segment should be such that the joint axis can be identified. If the joint angle remains constant, i.e. the links are rigidly connected, then $g_1(t) = R g_2(t)$, where R is the constant rotation matrix from the second to the first sensor frame. Therefore, (1) holds for any combination (j_1, j_2) , $j_1 = R j_2$, no matter what movement is performed. But it takes as little as rotating the first segment while the second is fixed, and then vice versa, to make the kinematic constraint become manifest in the gyroscopic measurement data.

The result of the Gauss-Newton algorithm may then, e.g., be used to calculate $(\hat{j}_1 \cdot g_1(t) - \hat{j}_2 \cdot g_2(t))$ which gives an estimate of the joint angle's time derivative. However, it should be noted that (1) is indifferent to the signs of both j_1 and j_2 . Therefore, we need a criterion to determine whether the gained estimates \hat{j}_1 and \hat{j}_2 are aligned or misaligned. In practice, this information can easily be gained by restricting the mounting of the sensor units such that, e.g., both z-axes globally point into the same half-space. However, if such information is not available, then alignment can as well be checked based on the measurement data, as we will see in Section III-C.




Fig. 1. Two rigid segments that are connected by a hinge joint, each of them equipped with a three-dimensional gyroscope (represented by its local coordinate system). The orientations of the gyroscopes towards their segments are assumed to be, and in many application actually are, unknown.




Fig. 2. Two rigid segments that are connected by a spheroidal joint, each of them equipped with a six-dimensional inertial measurement unit (IMU) consisting of a gyroscope and an accelerometer (represented by their common local coordinate system). Both the exact locations of the IMUs and their orientations towards their segments are assumed to be, and in many application actually are, unknown.

B. Constraints induced by spheroidal joints

Now we consider two links connected by a spheroidal joint, as depicted in Fig. 2. Since this joint has three degrees of freedom, there is no general relation between the measured angular velocities of the first and the second sensor. In order to exploit the kinematic constraints we need to incorporate the accelerometer readings. Let the accelerations of the sensors be $a_1(t)$ and $a_2(t)$ for the first and the second segment, respectively. And define o_1 and o_2 as the offset vectors from the joint center to the origin of the first and the second sensor frame, respectively². Then we claim that

$$\|a_1(t) - \Gamma_{g_1}(o_1)\|_2 - \|a_2(t) - \Gamma_{g_2}(o_2)\|_2 = 0 \quad \forall t, \quad (3)$$

$$\Gamma_{g_i}(o_i) := g_i(t) \times (g_i(t) \times o_i) + \dot{g}_i(t) \times o_i, \quad i = 1, 2,$$

where $\Gamma_{g_i}(o_i)$ is the (radial and tangential) acceleration due to rotation around the joint center. Therefore, $(a_1(t) - \Gamma_{g_1}(o_1))$ yields the acceleration of the joint center in the coordinates of the first local frame, which must equal $(a_2(t) - \Gamma_{g_2}(o_2))$ up to multiplication by some rotation matrix. Hence (3) holds for each instant in time and can be used to estimate o_1 and o_2 from a large number of data sets $(a_1(t_k), a_2(t_k), g_1(t_k), g_2(t_k))_{k=1}^N$. The gradients of the left-hand side terms is found to be

$$\frac{d(\|a_i(t) - \Gamma_{g_i}(o_i)\|_2)}{do_i} = \frac{\Gamma_{g_i}^T(a_i(t) - \Gamma_{g_i}(o_i))}{\|a_i(t) - \Gamma_{g_i}(o_i)\|_2}, \quad (4)$$

$$\Gamma_{g_i}^T(o_i) := (o_i \times g_i(t)) \times g_i(t) + o_i \times \dot{g}_i(t), \quad i = 1, 2,$$

where the matrix-multiplication representation of $\Gamma_{g_i}^T$ happens to be the transpose of the matrix-multiplication representation³ of the previously defined Γ_{g_i} .

Again, for a good estimation it is required that the motion is rich enough to make the constraint become manifest in the

²Note that, just as before, the subscripts 1 and 2 indicate that the vectors are in the coordinates of the first and second local frame, respectively.

³i.e. the matrix $G_{g_i} \in \mathbb{R}^{3 \times 3}$ that satisfies $G_{g_i} o_i = \Gamma_{g_i}(o_i) \quad \forall o_i \in \mathbb{R}^3$




Fig. 3. Kinematic simulation model of three segments connected by a hinge joint (axis marked blue) and a spheroidal joint. *Left*: Each simulated inertial measurement frame (red) is rigidly connected to its segment (black) with randomly chosen location and orientation. *Middle*: Random initial values are generated for the joint axis estimates of the two upper sensors (shifted to hinge joint center) and for the joint position estimates of the lower two sensors. *Right*: After a few seconds of leg and foot circling, and about ten Gauss-Newton steps, the estimates have converged to the true values. See also [9].

measurement data. Note, e.g., that if we restrict the (relative) motion of both segments to only one plane, then we basically have a hinge joint instead of a spheroidal joint, and (3) is true for all points along the (virtual) hinge joint axis. This as well as the feasibility of the estimation itself will be validated by simulation in the next section.

III. MODELING AND SIMULATION

A. A three-segments model

A kinematic simulation model is developed, which consists of three segments connected by a hinge and a spheroidal joint, see Fig. 3. Each simulated inertial measurement unit is rigidly connected to the respective segment. Prior to each simulation run, the positions and orientations of the sensors' coordinate systems are chosen randomly from reasonable intervals. The true joint axis and joint position, in the coordinates of the respective sensor frames, is calculated and stored for comparison with the estimates. Subsequently, the segments perform a user-defined motion including translation and rotation within the bounds of the kinematic constraints. Based on an adjustable sample rate, the accelerations and angular velocities of the sensors are computed in the coordinates of their local frames. A user-defined amount of measurement noise is added, and the data is then provided to the estimation algorithms.

B. Algorithm implementation

Assume that N data sets, precisely $(g_1(t_k), g_2(t_k))_{k=1}^N$, are provided, where $N \gg 4$. By restricting the joint axis estimates to unit length, the estimation problem becomes four-dimensional. Accordingly, \hat{j}_1 and \hat{j}_2 are parametrized

in spherical coordinates, i.e.

$$\begin{aligned} x &:= (\phi_1, \theta_1, \phi_2, \theta_2)^T, \\ \hat{j}_1 &= (\cos(\phi_1) \cos(\theta_1), \cos(\phi_1) \sin(\theta_1), \sin(\phi_1))^T, \\ \hat{j}_2 &= (\cos(\phi_2) \cos(\theta_2), \cos(\phi_2) \sin(\theta_2), \sin(\phi_2))^T, \end{aligned} \quad (5)$$

where ϕ_i and θ_i are inclination and azimuth of \hat{j}_i in the i^{th} sensor's coordinate system, $i = 1, 2$. We generate random initial values for both the inclinations and azimuths, see e.g. Fig. 3, and the following update procedure is applied:

- 1) Use (5) to calculate \hat{j}_1 and \hat{j}_2 from x .
- 2) Calculate the error vector $e \in \mathbb{R}^{N \times 1}$ defined by $e(k) := \|\hat{j}_1 \times g_1(t_k)\|_2 - \|\hat{j}_2 \times g_2(t_k)\|_2$, $k = 1, \dots, N$.
- 3) Use (2) and (5) to calculate the Jacobian $\frac{de}{dx}$ as well as its Moore-Penrose-pseudoinverse $\text{pinv}(\frac{de}{dx})$.
- 4) Update x by $x = x - \text{pinv}(\frac{de}{dx}) e$ and repeat from 1).

A similar update scheme is applied for the estimation of the spheroidal joint's offset vectors, but here x is the concatenation of o_1 and o_2 . The error is defined by $e(t) := \|a_1(t) - \Gamma_{g_1}(o_1)\|_2 - \|a_2(t) - \Gamma_{g_2}(o_2)\|_2$, and instead of (2) we use (4) for the calculation of the Jacobian. Therefore, and for the error equation itself, the time-derivatives of $g_1(t)$ and $g_2(t)$ are required. A noise-rejecting non-causal low-pass filter combined with a simple difference quotient approximation is used and will prove to give good results in Section IV, although more sophisticated approaches may exist.

It may seem that the proposed algorithms require large-scale numerics. However, the matrix that needs to be inverted in the calculation of the pseudoinverse is at most of dimension six. Furthermore, for normal-speed motions and at common sample rates of 50 Hz to 300 Hz, it is sufficient to use only every third to tenth sample, since data sets with

almost the same angular velocities hardly contribute to the estimation. This also holds the advantage that the proposed algorithms work well with considerably low sample rates.

In addition to the above, it should be mentioned that there are many other least-squares implementations that might be considered. As an online modification, e.g., one might start applying the given update procedure as soon as a minimum number of data sets, typically about ten, is available and then repeat the loop (multiple times) whenever a new data set is added. Then N increases gradually and every updated estimate serves as a (very good) initial condition to the next larger least-squares problem. Furthermore, please note that in both estimation algorithms one might also divide the difference of the two norms in the error equation by the mean of both, and thus use a relative error instead. However, in the presence of measurement errors we prefer to give more weight to the data sets that yield large norms. Hence we use the absolute errors, as defined above.

C. Simulation results

Both estimation algorithms are tested for various motions, sample rates, and noise amplitudes, both online and offline. It is found that two to three periods of (simultaneous) small-amplitude oscillations in the direction of each degree of freedom already yield enough data for an accurate estimation. Moreover, even for a few hundred randomly chosen initial conditions, the estimates always converge to the close proximity of the true values within five to ten iterations, see e.g. Fig. 3. At a signal-to-noise ratio of 100, the estimated joint axes and the true axis confine angles of less than 1° , while the joint position estimates differ from the true values by less than 3%.

Now we can use the hinge joint axis estimates to decompose, for each sample, the angular velocities into a rotation rate around the joint axis and a component orthogonal to that, i.e. the projection into joint plane.

$$g_{i,\text{proj}}(t) = g_i(t) - (g_i(t) \cdot j_i)j_i, \quad i = 1, 2 \quad (6)$$

Thereby we gain a little more insight into the idea behind (1), since, if we plot the projections for both sensors, as in Fig. 4, then we find that they have the same length for each sample. Moreover, if we corrected them point-wise by the respective rotation around the joint axis, then we would obtain two traces that are equal up to a rotation around the origin. This can easily be used to check the alignment of the axis estimates, since, if they are misaligned, then one of the traces will be mirrored.

Similarly, one might calculate $(a_i(t) - \Gamma_{g_i}(o_i))$, as defined in (3), for both sensors next to the spheroidal joint. Then for each sample one would obtain two vectors that are equal up to (three-dimensional) rotation around the origin by the (time-dependent) joint angles plus some constants.

For an illustration, please refer to [9].

As a final test, we apply the joint position estimation to the measurement data of the sensor units next to the hinge joint. It is found that almost every run converges to a different pair of offset vectors but each pair describes a point that happens to lie on the hinge joint axis. On the contrary, applying the joint axis estimation to the measurement data of the sensor units next to the spheroidal joint yields no convergence at all.




Fig. 4. The 3d measurements of two gyroscopes attached to the ends of a hinge joint are projected into the joint plane: For each moment in time the projections have the same length. Moreover, due to kinematic constraints the two curves are congruent up to a rotation around the origin by the (time-dependent) joint angle plus some constant. Since the joint angle only varies by $\pm 10^\circ$ in the given simulation, the curves are hardly distorted.

IV. EXPERIMENTAL RESULTS

For the sake of experimental validation of the proposed estimation algorithms, we apply them to measurement data from an IMU-based gait analysis. Wireless motion trackers⁴ are attached to the right thigh, shank and foot using elastic straps. The true position and orientation of the local sensor frames towards the ankle and knee are roughly⁵ determined using a combination of manual measurements and calibration movements, as in [5] and [2], respectively. Subsequently, the subject performs leg and foot circling for about five seconds and then walks about thirty meters at average speed. The accelerometer and gyroscope readings are recorded and provided to the algorithms developed above at a sample rate of 40 Hz.

One hundred runs are performed using random initial

⁴xsens MTw, <http://www.xsens.com>

⁵i.e. up to the limitations in accuracy that are inherent to these methods

values and different subsets⁶ of the available data to analyze both convergence and variance of the estimations. Results are presented in Fig. 5. In all runs, both the knee joint axis and the ankle joint position are identified correctly within less than twenty iterations. However, the error vector norms do not entirely converge to zero and the final values vary by about ± 0.01 . This is not surprising, since the sensor-to-leg connections are not very rigid, and since the knee and ankle are not exactly a hinge and a spheroidal joint. But apparently, the least-squares approach can cope with these inaccuracies and still yield accurate estimates. Of course, significantly higher accuracy and less variance is expected in more rigid setups such as robotic manipulators.




Fig. 5. Joint axis and position estimation from IMU-based gait analysis data. For one hundred random initial values and with different subsets of the available data stream under consideration, the estimates always converge to the true values (marked by semicircles) within less than twenty iterations. The final values show little variance despite the flexibilities in the mechanical setup.

V. CONCLUSIONS AND FUTURE WORKS

We proposed least-squares methods for the estimation of joint axes and positions from 6d inertial measurement data. Explicit analytic expressions of the required Jacobians have been provided as well as an example for algorithm implementation and a number of optional modifications. It was demonstrated that, regardless of the choice of initial estimates, convergence to the true values is obtained both in simulation and experiment. Due to the nature of the approach, very small sample rates suffice and neither integration nor knowledge on the sensor mounting is required. The drawback of potential misalignment of the

⁶Note, however, that all subsets include data from both the circling and the walking phase.

axis estimates has been addressed and two simple practical solutions have been suggested. Finally, the reliability of the algorithms in the context of human gait analysis has been demonstrated.

Further fields of application include robotic manipulators, linked vehicles, or any other mechanical setup in which rigid bodies are connected by joints. In any case, the sensor units can be attached in arbitrary position and orientation, and calibration movements become obsolete. The obtained estimates can be used to calculate joint angles, to transform the sensors readings into joint-related coordinate systems, or to apply bias-eliminating techniques as in [1]. Future work will be concerned with such extensions and with the identification of angle-dependent joint axes and positions.

VI. ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the German Federal Ministry of Education and Research (FKZ 01EZ1112), and Steffen Schäperkötter's help with the experiments, as well as Timo von Marcard's advice regarding the sensors.

REFERENCES

- [1] A.D. Young, "Use of Body Model Constraints to Improve Accuracy of Inertial Motion Capture", International Workshop on Wearable and Implantable Body Sensor Networks, Pages 180-186, International Conference on Body Sensor Networks, 2010.
- [2] K.J. ODonovan, R. Kamnik, D.T. O'Keefe, G.M. Lyons, "An inertial and magnetic sensor based technique for joint angle measurement", Journal of Biomechanics, Volume 40, Issue 12, 2007, Pages 2604-2611.
- [3] J. Favre, B.M. Jolles, R. Aissaoui, K. Aminian, "Ambulatory measurement of 3D knee joint angle, Journal of Biomechanics", Volume 41, Issue 5, 2008, Pages 1029-1035.
- [4] P. Cheng, B. Oelmann, "Joint-Angle Measurement Using Accelerometers and Gyroscopes: A Survey," Instrumentation and Measurement, IEEE Transactions on , Volume 59, Number 2, Pages 404-414, February 2010.
- [5] K. Liu, T. Liu, K. Shibata, Y. Inoue, "Ambulatory measurement and analysis of the lower limb 3D posture using wearable sensor system", International Conference on Mechatronics and Automation 2009, Pages 3065-3069.
- [6] R. Takeda, S. Tadano, A. Natorigawa, M. Todoh, S. Yoshinari, "Gait posture estimation using wearable acceleration and gyro sensors", Journal of Biomechanics, Volume 42, Issue 15, November 2009, Pages 2486-2494.
- [7] T. Liu, Y. Inoue, K. Shibata, "Development of a wearable sensor system for quantitative gait analysis", Measurement, Volume 42, 2009, Pages 978-988.
- [8] K. Liu, T. Liu, K. Shibata, Y. Inoue, R. Zheng, "Novel approach to ambulatory assessment of human segmental orientation on a wearable sensor system", Journal of Biomechanics, Volume 42, Issue 16, December 2009, Pages 2747-2752.
- [9] Supplemental material, e.g. 3d animations, is available at <http://www.control.tu-berlin.de/IMU-Based.Gait.Detection.and.Gait.Analysis>