Experimental evaluation of joint position estimation using inertial sensors

Fredrik Olsson

Department of Information Technology
Uppsala University, Uppsala, Sweden
Email: fredrik.olsson@it.uu.se

Kjartan Halvorsen

Department of Information Technology
Uppsala University, Uppsala, Sweden and
Department of Mechatronics
ITESM, Atizapan, Mexico
Email: kjartan.halvorsen@it.uu.se

Abstract—In this paper we study the problem of estimating the position of a joint that is connecting two rigid links in a biomechanical model. By equipping the two links with inertial sensors, which measure linear acceleration and angular velocity, it is possible to estimate the joint position. Estimation methods for this problem have been proposed before, but experimental evaluation and comparison between methods is lacking. The main contribution of this paper is an experimental evaluation of three different methods, a least-squares method and two iterative optimization methods that minimize a sum of squared errors and a sum of absolute errors. The iterative methods turns out to be superior to the least-squares method. Minimizing the sum of squared errors is faster but not as robust against outliers as minimizing the sum of absolute errors.

I. INTRODUCTION

Inertial sensors, specifically triaxial accelerometers and gyroscopes which measure linear acceleration and angular velocity in three dimensions, are widely available in our everyday lives. As microelectromechanical systems (MEMS), inertial sensors are staple components in today's smart phones and other wearable devices such as activity monitors, smart wristbands and watches. This has made inertial sensors attractive tools for analysing human motion, for example to perform full body motion capture [1], quantifying symptoms of diseases, tracking the progress of rehabilitation [2] or aiding elite sportsmen in their training [3]. In all of these applications, an important problem to consider is how the measurements from the sensors relate to the human body. A human body can be approximated as separate links that connect to each other through joints that allow the links to move in certain ways with respect to each other. Such a model of a human is what we refer to as a biomechanical model. Establishing the relationship between the inertial sensors and links and joints of the biomechanical model is referred to as calibrating the biomechanical model, and it is often a necessary step in the mentioned applications.

In this paper we consider the problem of estimating the position of a joint connecting two links in the biomechanical model. One inertial sensor will be attached to each link. We

assume the positions of the sensors to be fixed with respect to the center of the joint, but apart from this there will be no further prior knowledge assumed about the positions and orientations of the sensors. When the position of the joint center is fixed with respect to the reference frame of the sensors, it is possible to estimate the position of the joint using the measurements from the two inertial sensors. Similar problems have been studied previously in e.g. [4], [5], [6], but experimental evaluation and comparisons of the methods are lacking. Our main contribution with this paper is the experimental evaluation of three possible methods for estimating the joint position using the two inertial sensors. Evaluations were made using a rather simple experimental system, an artificial limb consisting of two rigid links connected by a spherical joint.

II. MODELING

A. Relating the inertial measurements to the biomechanical model

Let L_1 and L_2 denote local (Cartesian) reference frames that are fixed in sensor 1 and 2, respectively, and G denote a global earth-fixed reference frame. The measured accelerations a relate to the measured angular velocities ω and the angular accelerations $\dot{\omega}$ in the following way

$$a_1 = a_0^{L_1} + \omega_1 \times (\omega_1 \times r_1) + \dot{\omega}_1 \times r_1 \quad (1)$$

$$a_2 = a_0^{L_2} + \omega_2 \times (\omega_2 \times r_2) + \dot{\omega}_2 \times r_2, \quad (2)$$

where a_0 is the linear acceleration of the center of rotation expressed in L_1 and L_2 , respectively, as indicated by the superscripts. The measured accelerations, angular velocities and angular accelerations, are expressed in the local reference frame associated with the subscripts. A similar notation is used for the position of the joint, r , when expressed in the local reference frames. The described model is illustrated in Figure 1. We can rewrite (1)-(2) as

$$a_1 = a_0^{L_1} + K(\omega_1, \dot{\omega}_1)r_1 \quad (3)$$

$$a_2 = a_0^{L_2} + K(\omega_2, \dot{\omega}_2)r_2, \quad (4)$$

Fig. 1. The biomechanical model and inertial sensors, where the joint positions with respect to the sensors, r_1 and r_2 are to be estimated.

where

$$K(\omega, \dot{\omega}) = \begin{bmatrix} -\omega_y^2 - \omega_z^2 & \omega_x \omega_y - \dot{\omega}_z & \omega_x \omega_z + \dot{\omega}_y \\ \omega_x \omega_y + \dot{\omega}_z & -\omega_x^2 - \omega_z^2 & \omega_y \omega_z - \dot{\omega}_x \\ \omega_x \omega_z - \dot{\omega}_y & \omega_y \omega_z + \dot{\omega}_x & -\omega_x^2 - \omega_y^2 \end{bmatrix}, \quad (5)$$

where the x, y, z subscripts denote the coordinates of the three dimensional vectors. To simplify notation we will use $K(\omega_1, \dot{\omega}_1) = K_1$ and $K(\omega_2, \dot{\omega}_2) = K_2$.

B. Measurement models

The accelerometers and gyroscopes provide us with noisy measurements of the accelerations and the angular velocities. Simple models of the measurements at sample instant k are given by

$$y_a(k) = a(k) + e_{a,k}, \quad e_{a,k} \sim \mathcal{N}(0, \Sigma_a) \quad (6)$$

$$y_\omega(k) = \omega(k) + e_{\omega,k}, \quad e_{\omega,k} \sim \mathcal{N}(0, \Sigma_\omega), \quad (7)$$

where we say that the measurement is equal to the physical quantity plus some additive Gaussian measurement noise. Inertial sensors, especially lower quality sensors, tend to also have a time varying bias, causing the measurements to slowly drift over time. A simple way to estimate the bias is to collect data from stationary sensors and see how much they differ from the expected stationary accelerations and angular velocities. This initially estimated bias can then be assumed to be stationary over short time periods and removed from the data. Angular accelerations are also needed by the methods, and these must

be estimated from the gyroscope measurements. We will use a simple second-order symmetric difference quotient

$$\dot{\omega}(k) \approx y_{\dot{\omega}}(k) = \frac{y_\omega(k+1) - y_\omega(k-1)}{2h}, \quad (8)$$

where h is the sample period. Other, higher order differentiation techniques may also be used. When estimating the position of the joint we will make use of these noisy measurements, so when we refer to a , ω and $\dot{\omega}$ in the upcoming sections we will mean the measurements given by (6)-(8) and not the true physical quantities.

III. ESTIMATION

A. Least-squares estimate

Note that (3)-(4) is linear in a_0 and $r = [r_1 \ r_2]^\top$. Let $R_{1,2}$ denote the rotation matrices which rotates vectors expressed in the $L_{1,2}$ frames to the G frame. For example we have that

$$R_1 a_0^{L_1} = a_0^G.$$

By using these rotation matrices we can express (3)-(4) in the G frame

$$\begin{aligned} R_1 a_1 - R_2 a_2 &= a_0^G - a_0^G + R_1 K_1 r_1 - R_2 K_2 r_2 \\ &= R_1 K_1 r_1 - R_2 K_2 r_2, \end{aligned} \quad (9)$$

which is a system of equations which is linear in r . Once both accelerations are expressed in a common reference frame, a_0 is eliminated from the equation by subtracting one acceleration from the other. With a set of N available measurements, the following system of equations can be constructed

$$\tilde{K} r = \tilde{a} \quad (10)$$

$$\tilde{K} = \begin{bmatrix} \tilde{K}(1) \\ \vdots \\ \tilde{K}(N) \end{bmatrix} = \begin{bmatrix} R_1(1)K_1(1) & -R_2(1)K_2(1) \\ \vdots & \vdots \\ R_1(N)K_1(N) & -R_2(N)K_2(N) \end{bmatrix} \quad (11)$$

$$\tilde{a} = \begin{bmatrix} \tilde{a}(1) \\ \vdots \\ \tilde{a}(N) \end{bmatrix} = \begin{bmatrix} R_1(1)a_1(1) - R_2(1)a_2(1) \\ \vdots \\ R_1(N)a_1(N) - R_2(N)a_2(N) \end{bmatrix}, \quad (12)$$

where $\tilde{K} \in \mathbb{R}^{3N \times 6}$ and $\tilde{a} \in \mathbb{R}^{3N \times 1}$. A least-squares (LS) estimate of r can be found by minimizing a cost function which corresponds to the sum of squared errors of the linear equations

$$\hat{r}_{LS} = \arg \min_r \sum_{k=1}^N \left\| \tilde{K}(k)r - \tilde{a}(k) \right\|_2^2, \quad (13)$$

which has the closed form solution

$$\hat{r}_{LS} = \tilde{K}^\dagger \tilde{a}, \quad (14)$$

where $\dagger$ denotes the Moore-Penrose pseudo inverse.

The rotation matrices $R_{1,2}$ can be found by estimating the orientation of the inertial sensors with respect to G . There are several methods that one can use to estimate the orientations of the inertial sensors [7], and some inertial sensors include

internally computed orientation estimates. However, since estimates of $R_{1,2}$ are needed to compute $\hat{r}_{LS}$ this introduces an additional source of uncertainty that can be hard to quantify. It should be noted that the LS approach presented here extends the methods presented in [5] and [6]. These methods assume that $a_0 = 0$, which would allow for separate estimation of r_1 and r_2 without having to know the orientation of the sensors. There are not many scenarios where the assumption that the center of the joint is at rest holds when studying human motion, however. So it is less restrictive and more practical to consider the case when $a_0 \neq 0$.

B. Estimation using iterative optimization methods

Another way of estimating r is to use iterative optimization methods. The method proposed by Seel et al. in [4] minimizes a sum of squared errors (SS) cost function

$$\hat{r}_{SS} = \arg \min_r \sum_{k=1}^N e(k, r)^2 \quad (15)$$

$$e(k, r) = \|a_1(k) - K_1(k)r_1\|_2 - \|a_2(k) - K_2(k)r_2\|_2, \quad (16)$$

where $\|\cdot\|_2$ stands for the Euclidean norm. A nice property of defining the errors $e(k, r)$ as the difference between two Euclidean norms is that the Euclidean norm is invariant to rotations. Therefore, unlike the LS method, we do not need to estimate the orientation of the sensors to remove a_0 from the equation since $\|a_0^{L_1}\|_2 = \|a_0^{L_2}\|_2$. The optimization problem (15) can be solved efficiently using a Gauss-Newton method [8] as shown in [4], which makes use of the partial derivatives

$$\frac{\partial e(k, r)}{\partial r_1} = -\frac{K_1(k)^\top e_1(k, r_1)}{\|e_1(k, r_1)\|_2} \quad (17)$$

$$\frac{\partial e(k, r)}{\partial r_2} = \frac{K_2(k)^\top e_2(k, r_2)}{\|e_2(k, r_2)\|_2}, \quad (18)$$

where

$$e_i(k, r_i) = a_i(k) - K_i(k)r_i, \quad i = 1, 2. \quad (19)$$

An alternative to minimizing the sum of squared errors is to instead minimize the sum of absolute errors (SA)

$$\hat{r}_{SA} = \arg \min_r \sum_{k=1}^N |e(k, r)|. \quad (20)$$

A well known property of this type of cost function is robustness against outliers in the data since large residuals are weighted by their absolute value, not by their squared value in the cost function. By outliers we mean measurements that contain some sort of disturbance that the model does not take into account. One common disturbance comes from the fact that human body segments are non-rigid. The sensors are placed on soft tissue, sometimes with clothing between the sensors and the skin. This violates the assumption that the sensor position with respect to the joint is fixed at all times since the surface on which the sensors are placed can move to some degree. These types of disturbances are

known as soft tissue artifacts (STA) [9]. Since the optimization problem given by (20) is not a sum of squares minimization problem, it is not possible to use the Gauss-Newton method to solve this optimization problem. Other gradient descent-type methods [10] may be used, which require the partial derivatives (subgradients)

$$\frac{\partial |e(k, r)|}{\partial r_1} = -\text{sgn}(e(k, r)) \frac{K_1(k)^\top e_1(k, r_1)}{\|e_1(k, r_1)\|_2} \quad (21)$$

$$\frac{\partial |e(k, r)|}{\partial r_2} = \text{sgn}(e(k, r)) \frac{K_2(k)^\top e_2(k, r_2)}{\|e_2(k, r_2)\|_2}, \quad (22)$$

where

$$\text{sgn}(x) = \begin{cases} -1 & x < 0 \\ 0 & x = 0 \\ 1 & x > 0 \end{cases}. \quad (23)$$

It is important to note that neither of the estimation methods described in this section guarantee that a unique solution exists in general. If the data contains unmodeled disturbances or outliers, multiple solutions or local minima may exist. The resulting estimates may then in some cases describe the outliers rather than the kinematics of the model of interest.

IV. EXPERIMENT

An experiment was performed using a system consisting of two rigid cylindrical links joined together by a spherical joint, which allows the links to rotate with up to two degrees of freedom with respect to each other. Two Xsens MTw Awinda wireless inertial sensors [11] were placed near the ends of the links. The setup of the experimental system can be seen in Figure 2. The links are rigid such that the position of the joint relative to the sensors can be considered fixed for the full duration of the experiment. The positions of the joint relative to the sensors were measured beforehand to be

$$r_1 = [334 \quad 0 \quad 28]^\top (\pm 2) \quad (24)$$

$$r_2 = [-325 \quad 0 \quad 28]^\top (\pm 2), \quad (25)$$

expressed in the $L_{1,2}$ frames in millimeters. The experiment consisted of collecting inertial measurements while picking up the system and performing certain motions by hand. The sample rates of the sensors were 100 Hz and a total of 15219 data samples were collected during the experiment. The data can be divided into segments $S_i, i = 1, 2, \dots, 8$, each of which contains one specific type of motion, see Figure 3. The motions contained in each segment are:

- S_1 : The system was at rest on a flat surface. This segment was used to calibrate gyroscope bias and noise covariance. Contains 1354 samples.
- S_2 : The system was picked up and held such that the two segments were in an approximately vertical orientation. Contains 1132 samples.
- S_3 : The links were moved back and forth with respect to the joint. The majority of the rotation aligned with the y -axes of the sensors. Contains 1625 samples.

Fig. 2. The system that was used in the experiment with two Xsens MTw inertial sensors attached.

Fig. 3. Data collected from sensor 1 during the experiment.

- S_4 : The links were moved back and forth with respect to the joint. The majority of the rotation aligned with the z -axes of the sensors. Contains 1372 samples.
- S_5 : The links were moved back and forth with respect to the joint. Alternated rotations around the y - and z -axes of the sensors. Contains 2213 samples.
- S_6 : The links were spun such that the bottom link followed an approximately clockwise circular trajectory. The system was flipped upside down in the middle of the segment and the motion repeated. Contains 2666 samples.
- S_7 : Same as S_6 but counter-clockwise motion. Contains 2698 samples.
- S_8 : The system was placed back on the flat surface as in S_1 . Contains 2159 samples.

V. RESULTS

A. Metrics for evaluation of estimators

The performance of the estimators are evaluated in terms of the average Euclidean error and its standard deviation

$$E_{\text{avg}}(\hat{r}_i) = \frac{1}{M} \sum_{j=1}^M E_j(\hat{r}_i), \quad i = 1, 2 \quad (26)$$

$$\text{Std}[E_j(\hat{r}_i)] = \sqrt{\frac{1}{M} \sum_{j=1}^M (E_j(\hat{r}_i) - E_{\text{avg}}(\hat{r}_i))^2} \quad (27)$$

$$E_j(\hat{r}_i) = \|\hat{r}_i(j) - r_i\|_2, \quad (28)$$

TABLE I
SHOWS E_{AVG} [MM] AND STD $[E_j]$ [MM] OVER $M = 100$ RUNS WITH $N = 500$ SIMULATED MEASUREMENTS.

Data	Method	$\hat{r}_1$		$\hat{r}_2$	
		E_{avg}	Std $[E_j]$	E_{avg}	Std $[E_j]$
Simulated	LS	0.22	0.09	0.23	0.10
	SS	0.07	0.03	0.07	0.03
	SA	0.08	0.03	0.09	0.04

where $E_j(\hat{r}_i)$ denotes the Euclidean norm of the error for the j :th estimate of r_i . The average Euclidean error $E_{\text{avg}}(\hat{r}_i)$ and its standard deviation $\text{Std}[E_j(\hat{r}_i)]$ are computed for a total of M estimates. The rationale for choosing the average Euclidean error over e.g. the more commonly used root mean square error (RMSE) is that it may be interpreted as the average physical distance between estimates with respect to the true joint center [12]. This results in a more intuitive understanding as to how close the estimators are to estimating the true position of the joint center. The standard deviation of the Euclidean error give us information about the variation or consistency of the estimates.

B. Method verification using simulated data with ideal excitation

To see that the methods were correctly implemented and working we first used simulated measurements for estimating $r_{1,2}$. Angular velocities and angular accelerations were randomly generated from zero-mean Gaussian distributions with standard deviations of 10 rad/s and 10 rad/s², respectively. Accelerometer measurements were then generated from the models (1) and (2) using the measured sensor positions (24) and (25). The accelerations of the joint a_0 were uniformly random rotations of the gravitational acceleration vector $g = [0 \ 0 \ 9.82]^\top$. Additive zero-mean Gaussian noise with variances $10^{-2} \text{ m}^2/\text{s}^4$, $10^{-5} \text{ rad}^2/\text{s}^2$ and $5 \times 10^{-2} \text{ rad}^2/\text{s}^4$, were added to the accelerations, angular velocities and angular accelerations, respectively. The first two noise levels are typical for the accelerometers and gyroscopes found in the Xsens MTw sensors, and the noise added to the angular accelerations takes into account the second order finite difference quotient that is used to differentiate the noisy gyroscope signals. In addition to measurement noise, orientation errors were simulated as zero-mean Gaussian random variables with a standard deviation of 5° . By allowing the angular velocities and angular accelerations to be randomly generated like this, we do not take constraints of the real biomechanical model into account, such as the physically possible range of motion. The randomly generated measurements give us an indication of how well the methods would perform if the system could be ideally excited, with no constraints limiting the motion. A total of $M = 100$ estimation runs were performed using $N = 500$ randomly generated measurements in each run. The results for these runs are shown in Table I.

TABLE II
SHOWS E_{AVG} [MM] AND STD $[E_j]$ [MM] OVER $M = 100$ RUNS WITH
 $N = 6280$ SAMPLES FROM SEGMENTS S_3 - S_7 .

Data	Method	$\hat{r}_1$		$\hat{r}_2$	
		E_{avg}	Std $[E_j]$	E_{avg}	Std $[E_j]$
S_3 - S_7	LS	16.8	0.0	11.3	0.0
	SS	13.8	0.0	8.9	0.0
	SA	13.0	0.0	7.1	0.0

C. Estimation using the experimental data

The joint positions $r_{1,2}$ were estimated using the data collected from the experiment described in Section IV. At first, estimation was done using data from segments S_3 to S_7 . To avoid using non-informative stationary data, a threshold of 1.2 rad/s was used to exclude all samples where the norm of the gyroscope measurement was below this threshold. This threshold was also used in all subsequent estimations. As a result, a total of $N = 6280$ measurements were selected from the segments S_3 to S_7 . Using the same set of measurements, $M = 100$ estimation runs were performed with the same initial conditions for the two iterative estimates $\hat{r}_{SS}$ and $\hat{r}_{SA}$. Initial conditions were randomly generated for each estimation run from a standard zero-mean Gaussian distribution with unit-variance. The rotation matrices $R_{1,2}$ required for the LS estimate, $\hat{r}_{LS}$ were estimated using an extended Kalman filter (EKF) with orientation deviation states [7]. The global reference frame G was the geodetic reference frame with the x -axis pointing north, the y -axis pointing west and the z -axis pointing up. The results for these runs are shown in Table II.

To see how the accuracy of the estimates may vary for different types of motion, the joint positions were estimated using samples from each segment separately. From each segment $N = 500$ measurements were randomly chosen for each estimation run for a total of $M = 100$ runs per segment. For segments S_6 and S_7 , 250 measurements were randomly chosen from the first and the second half of the segments, respectively, since the sensors were flipped upside down in the middle of the segment. This was to make sure we always got measurements of the rotation of both sensors. The results for these runs are shown in Table III.

D. Estimation with simulated soft tissue artifacts

In this experiment we investigated how the methods performed when the inertial measurements were corrupted by disturbances. Possible disturbances when analysing real human motion may appear in the form of STAs. Since the links in the experimental system are rigid we simulated these disturbances. A model that has been shown to emulate STAs in the thigh reasonably well expresses the displacement of markers attached to the skin as a linear combination of the hip joint angles [13]. Differentiating this model twice w.r.t. time yields a linear relationship between the linear acceleration of the displacement of skin and the angular acceleration. Using this model structure as a basis, we generated STAs as an

TABLE III
SHOWS E_{AVG} [MM] AND STD $[E_j]$ [MM] OVER $M = 100$ RUNS WITH
 $N = 500$ RANDOMLY PICKED SAMPLES FROM DIFFERENT SEGMENTS.

Data	Method	$\hat{r}_1$		$\hat{r}_2$	
		E_{avg}	Std $[E_j]$	E_{avg}	Std $[E_j]$
S_3	LS	332.7	19.4	181.8	18.3
	SS	17.5	0.5	6.6	0.4
	SA	17.5	2.3	6.8	1.6
S_4	LS	806.5	13.4	655.5	11.5
	SS	15.9	0.4	5.1	0.3
	SA	15.0	1.3	6.0	0.7
S_5	LS	325.6	23.4	129.3	27.5
	SS	13.8	0.5	6.8	0.5
	SA	13.0	1.3	7.3	1.0
S_6	LS	40.0	6.4	33.9	10.5
	SS	24.2	23.6	10.7	8.8
	SA	13.1	0.2	7.4	0.3
S_7	LS	20.9	8.1	21.1	7.1
	SS	22.4	12.7	9.4	2.4
	SA	13.0	0.3	7.1	0.3

TABLE IV
SHOWS E_{AVG} [MM] AND STD $[E_j]$ [MM] OVER $M = 100$ RUNS WITH
 $N = 6280$ SAMPLES FROM SEGMENTS S_3 - S_7 . WITH SIMULATED
DISTURBANCES ADDED TO THE ACCELEROMETER MEASUREMENTS.

Data	Method	$\hat{r}_1$		$\hat{r}_2$	
		E_{avg}	Std $[E_j]$	E_{avg}	Std $[E_j]$
S_3 - S_7	LS	31.6	12.2	22.7	12.4
	SS	17.6	5.6	11.2	4.2
	SA	15.2	3.7	10.4	4.0

additive disturbance to the acceleration measurements as a linear combination of the angular acceleration measurements

$$d_a = H\dot{\omega}, \quad (29)$$

where H is a 3×3 matrix containing the STA model parameters. In practice, the elements of H depend on multiple factors such as sensor placement and the physiology of the person. Therefore, the elements of H were randomized from a zero-mean Gaussian distribution with a standard deviation of $\frac{18}{\pi}$ mm/rad at the start of each estimation run. This generated H matrices with elements of a magnitude comparable to what was found in [13]. A total of $M = 100$ estimation runs were performed using the same $N = 6280$ samples from segments S_3 to S_7 that were used to produce the results shown in Table II. The results for these runs are shown in Table IV.

E. Analysis of the cost functions of the iterative method

Throughout the estimation runs we found that the iterative methods performed significantly worse than expected for certain parts of the data, see Table III. We analyzed how the cost functions (15) and (20) depend on the values of the x -coordinates of $\hat{r}_1$ and $\hat{r}_2$ in a small neighborhood around the true joint positions r_1 and r_2 . Figure 4 show the level curves of the cost functions for a specific set of data from

one estimation run using 250 random samples from the first half and 250 random samples from the second half of the S_6 segment. Similarly, Figure 5 show the level curves of the cost functions, but when using 500 random samples from the first half of the S_6 segment instead.

From Figure 4 we see that the SS cost function has two distinct local minima for this specific set of data, while the SA cost function only has one global minima. This would explain why the standard deviations of the SS estimates were significantly larger compared to the SA estimates for segments S_6 and S_7 . It is clear from this example that some of the random data samples picked during the estimation runs can cause multiple local minima to appear. The Gauss-Newton method will then converge to the local minima closest to the initial estimate.

Figure 5 show us what happens when the majority of the data consists of rotations of only one of the sensors. We see that the cost functions are much steeper in the r_2 direction than the r_1 direction since the first half of the S_6 segment consists mostly of rotations of the second sensor. A problem that may happen in cases such as these is that the iterative methods may converge to an r_1 coordinate that is anywhere in the valley, since the level curves of the cost function are nearly flat in one direction. This shows how important it is to choose data where both sensors have been sufficiently excited by the motion.

VI. DISCUSSION

A. Motion is important

The functionality of the three estimators was verified by the results shown in Table I, where measurements were simulated with an ideal excitation. However, a motion that would allow for such excitation is impossible to perform under the conditions presented by our experimental setup and in most scenarios involving human motion due to limitations in mobility. The experimental data we collected contains multiple different types of motion, and even when using data from segments S_3 - S_7 , the average Euclidean error increases for all three estimators as shown by the results in Table II when compared to Table I. All three estimators were consistent when using the data from the S_3 - S_7 segments however, as is evident from the standard deviations, which were smaller than 0.05 mm.

The importance of the motion becomes clearer from the results shown in Table III, where data was randomly picked from single segments. The errors vary significantly across the five segments. The LS estimator in particular show poor performance for segments S_3 - S_5 . The motion during those segments is mostly constrained to one or two planes which makes the K -matrix from (5) close to singular, causing numerical issues when computing the LS estimate. In Section V-E we also showed the importance of having sufficient rotation of both sensors as the cost functions tend to become flat in some direction otherwise, as shown in Figure 5. Therefore, the best strategy is to combine multiple types of motion which allows rotation of both sensors.

Fig. 4. The level curves of the SS and SA cost functions in a neighborhood around the true x -coordinates. Using $N = 500$ random samples from both halves of segment S_6 .

B. The main difference of the estimators lies in the cost functions

The three estimators minimize different cost functions to find an estimate of r . The cost functions are similar in the sense that they consist of sums of errors that should go to zero as the relationships in (1)-(2) become satisfied. However, our measurements are corrupted by noise and possibly other disturbances that are not taken into account by the model. Measurements that deviate significantly from the model, so called outliers, could have a significant impact on the estimates depending on how the errors are defined in the cost functions.

The LS (13) and the SS (15) cost functions model the errors differently. The LS cost function have errors that are linear in the unknown parameter r , but requires additional information about the orientation of the sensors, while the SS cost function have errors that are nonlinear in r , but are invariant to the orientation. What the LS and SS estimators have in common is that they seek to minimize the sum over squared errors, which

Fig. 5. The level curves of the SS and SA cost functions in a neighborhood around the true x -coordinates. Using $N = 500$ random samples from only the first half of segment S_6 .

gives these estimators some nice numerical properties. The LS estimates have a closed form solution, which can be computed efficiently. The SS estimates can be found using a Gauss-Newton method, which typically converges faster than gradient descent methods, since information about the Hessian of the cost function is utilized. However, since the cost functions contain the squared errors, it also means that large errors can have a significant impact on the estimates. Outliers can give rise to multiple local minima as in Figure 4, and may pull the estimates away from the true quantity, which will be discussed further in Section VI-D.

The SA cost function (20) models the errors similarly to the SS cost function, but contains a sum of the absolute errors rather than the squared errors. Outliers will therefore have a less significant impact on the value of the cost function, and will thus affect the estimates less. However, the SA cost function requires more effort to find its minimum compared to the SS cost function. The Hessian of the SA cost function

cannot be formed since the absolute value is not differentiable in the origin, and we have to rely solely on the information of the subgradients, resulting in slower convergence.

C. The least-squares estimator relies on accurate orientation estimates

Another important distinction of the LS estimator is that it requires orientation estimates in the form of the rotation matrices $R_{1,2}$. An EKF with orientation deviation states [7] was used to obtain orientation estimates for the whole data set. In essence, this type of EKF uses the gyroscope measurements to track the dynamic rotation of an inertial sensor. Since this involves integrating noisy measurements, the orientation estimates will have an error that increases in magnitude over time. To compensate for this, the EKF may also incorporate measurements from an accelerometer and/or a magnetometer. A stationary accelerometer measures the local gravitational acceleration, which is an observation of the sensor's inclination and a magnetometer measures the Earth magnetic field which is an observation of the heading. However, to incorporate these measurements in the EKF one has to make some assumptions.

The accelerometer is assumed to be stationary, which is clearly not the case for most of our data. Therefore, only accelerometer measurements of magnitude within $\pm 5\%$ of the magnitude of the local gravitational acceleration is used as an observation of inclination in the EKF. Choosing this threshold is important. A small threshold may cause the integration error to grow, since there will be fewer inclination observations available. On the other hand, a large threshold include more erroneous inclination observations since the assumption of a stationary sensor is violated more. We found 5% to be a good threshold for this data set. However, while tuning this threshold we found that the LS estimation errors varied significantly, which indicates that the LS estimator is sensitive to the orientation estimates. A good strategy could be to record a motion that allows the sensor to be close to stationary at regular intervals. A smaller threshold could then be used since the data set is designed to include regular inclination observations, while excluding most data that violates the assumption of a stationary sensor.

The Xsens MTw sensors used also include magnetometers which could be used to obtain heading observations. However, magnetometers measure magnetic disturbances in addition to the Earth magnetic field, which makes them unreliable in certain environments. We found that only using magnetometer measurements from the S_1 segment to obtain an initial heading of the sensors, but not using the magnetometer measurements as heading observations in the EKF, worked best for our data set.

If we represent the unknown orientation estimation errors as rotation matrices $E_{1,2}$ equation (9) is not satisfied. We get

$$\begin{aligned} E_1 R_1 a_1 - E_2 R_2 a_2 &= \\ &= (E_1 - E_2) a_0^G + E_1 R_1 K_1 r_1 - E_2 R_2 K_2 r_2 \\ &\neq E_1 R_1 K_1 r_1 - E_2 R_2 K_2 r_2, \end{aligned} \quad (30)$$

where the unknown acceleration of the joint center a_0 remains in the equation. Large orientation estimation errors in combination with large accelerations of the joint center are therefore detrimental to the LS estimator. As mentioned in Section V-C, we only used measurements where the magnitude of the gyroscope measurement was larger than 1.2 rad/s. This threshold also turned out to be important for the LS estimator. It is clear from (30) that the $K_{1,2}$ matrices carry information about $r_{1,2}$. Since $K_{1,2}$ depend on the angular velocities and angular accelerations, choosing to only use measurements of sufficiently large magnitude may suppress the error term, $(E_1 - E_2)a_0^G$ in (30).

The additional work needed to tune an orientation estimation method to work well with a particular data set may be seen as a disadvantage of the LS estimator when compared to the SS and SA estimators. However, it is possible that a better strategy for tuning the orientation estimation and selecting a threshold for the magnitude of the measurements could improve the performance of the LS estimator.

D. The SA estimator is least sensitive to outliers

The results in Table III shows that for the segments S_6 and S_7 the SS estimator performs significantly worse than the SA estimator. The motion in those segments consisted of spinning the links in an approximately circular trajectory. The spherical joint in our experimental setup is surrounded by cylindrical metallic edges, that collided occasionally during the recording of these motions. The impact from these collisions propagated to the end of the links and were picked up by the sensors, which can be seen as some of the distinct impulse like peaks in Figure 3. As the model given by (1)-(2) describes the acceleration of the sensors as the acceleration of the joint center plus a rotation of the links around the joint center, it does not model these impulses. We can therefore conclude that our data set contain some outliers, particularly located in the S_6 and S_7 segments. It then makes sense why the SA estimator performed better for those segments, as previously discussed in Section VI-B the SA cost function gives less weight to large errors than the SS cost function.

In human motion analysis, the measurements will contain different types of outliers, such as STAs caused by the movement of the skin with respect to the bone. We simulated STAs as an additive disturbance to the accelerometer measurements, comparable to the thigh STA model introduced in [13]. Comparing the results in Table IV with Table II we see that all three estimators perform worse with these simulated STAs. The SA estimator seem to perform slightly better than the SS estimator. However, taking into account the uncertainty in measuring the position of the joint center, the difference in the accuracy of the SS and SA estimator lies within the margin of the measurement uncertainty. Therefore, we conclude that the SS and SA estimators perform equally well when using the $N = 6280$ samples from segments $S_3 - S_7$, both with and without simulated STAs.

VII. CONCLUSION

The main contribution of this paper has been the experimental evaluation of three different methods for estimating the joint position. The estimation is done using inertial measurements collected from sensors attached to the two links that connect via the joint. In terms of accuracy, the SS and SA estimators performed better over all than the LS estimator. The SS estimator was shown to be more sensitive to impulse like outliers created by colliding metallic parts of the experimental setup, while the SA estimator performed well under these conditions. However, the SS estimator have the advantage of faster convergence. The LS estimator was found to be most sensitive to the type of motion recorded and to be highly dependent on accurate orientation estimates, which makes it less reliable than the SS and SA estimators. Depending on what requirements a potential user of these estimators have in terms of robustness and computation time, either the SS or SA estimators may be used.

ACKNOWLEDGMENT

Thanks to Dave Zachariah and Andreas Svensson for proof-reading the manuscript and for providing useful feedback that contributed to improving the quality of the paper.

REFERENCES

- [1] D. Roetenberg, H. Luinge, and P. Slycke, "Xsens MVN: full 6DOF human motion tracking using miniature inertial sensors," Xsens Technologies, Tech. Rep., April 2013.
- [2] S. Patel, H. Park, P. Bonato, L. Chan, and M. Rodgers, "A review of wearable sensors and systems with applications in rehabilitation," *Journal of NeuroEngineering and Rehabilitation*, vol. 9, no. 1, April 2012.
- [3] D. A. James, *The Application of Inertial Sensors in Elite Sports Monitoring*. New York, NY: Springer New York, 2006, pp. 289–294.
- [4] T. Seel, T. Schauer, and J. Raisch, "Joint axis and position estimation from inertial measurement data by exploiting kinematic constraints," in *IEEE International Conference on Control Applications*, 2012, pp. 45–49.
- [5] R. S. McGinnis and N. C. Perkins, "Inertial sensor based method for identifying spherical joint center of rotation," *Journal of Biomechanics*, vol. 46, pp. 2546–2549, 2013.
- [6] M. Crabolu, D. Pani, and A. Cereatti, "Evaluation of the accuracy in the determination of the center of rotation by magneto-inertial sensors," in *Sensors Applications Symposium IEEE*, 2016.
- [7] M. Kok, "Probabilistic modeling for sensor fusion with inertial measurements," Ph.D. dissertation, Linköping University, 2016.
- [8] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. Springer Series in Operations Research, 2006.
- [9] A. Leardini, L. Chiari, U. D. Croce, and A. Cappozzo, "Human movement analysis using stereophotogrammetry: Part 3. soft tissue artifact assessment and compensation," *Gait and Posture*, vol. 21.2, pp. 212–225, 2005.
- [10] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [11] Xsens, "MTw Awinda Wireless Motion Tracker," Online: <https://www.xsens.com/download/pdf/documentation/mtw2-awinda/MTw2-Awinda.pdf>, accessed February 2017.
- [12] X. R. Li and Z. Zhao, "Evaluation of estimation algorithms part i: Incomprehensive measures of performance," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 4, pp. 1340–1358, 2006.
- [13] V. Camomilla, A. Cereatti, L. Chéze, and A. Cappozzo, "A hip joint kinematic driven model for the generation of realistic thigh soft tissue artefacts," *Journal of Biomechanics*, vol. 46, no. 3, pp. 625–630, 2013.