

Probabilistic Stand Still Detection using Foot Mounted IMU

Jonas Callmer, David Törnqvist and Fredrik Gustafsson

Division of Automatic Control

Linköping University

Linköping, Sweden

[callmer](mailto:callmer@isy.liu.se), [tornqvist](mailto:tornqvist@isy.liu.se), [fredrik](mailto:fredrik@isy.liu.se)@isy.liu.se

Abstract — We consider stand still detection for indoor localization based on observations from a foot-mounted inertial measurement unit (IMU). The main contribution is a statistical framework for stand-still detection, which is a fundamental step in zero velocity update (ZUPT) to reduce the drift from cubic to linear in time. First, the observations are transformed to a test statistic having non-central chi-square distribution during zero velocity. Second, a hidden Markov model is used to describe the mode switching between stand still, walking, running, crawling and other possible movements. The resulting algorithm computes the probability of being in each mode, and it is easily extendable to a dynamic navigation framework where map information can be included. Results of first mode probability estimation, second map matching without ZUPT and third step length estimation with ZUPT are provided.

Keywords: Indoor localization, stand still detection, HMM, ZUPT

1 Introduction

The problem of indoor localization has received an increasing amount of attention in the last couple of years [1, 2, 3, 5, 4, 8]. The desire to accurately track the position of first responders or military personnel, or to provide a positioning aid for civilians in shopping malls and airports, has led to a transition from robot sensor platforms to human ones.

To track a person, a variety of sensors can be used. A foot or body mounted IMU with accelerometers, gyros and magnetometers is simple and cheap and is therefore a very common sensor. It is usually supported by a range measuring radio device such as WiFi or UWB [8] or is fused with preexisting maps for enhanced tracking precision [8].

The IMUs used for indoor localization are small and cheap and consequently perform quite poorly. There is commonly a drift in the gyros causing the orientation




Figure 1: Localization experiment without ZUPT. Red dots are the position hypotheses and the blue dot and line is the mean position and the heading measurement, respectively. The lack of ZUPT means that we have a large step length uncertainty, which causes the hypotheses to spread all along the corridor.

estimate to be incorrect. Since the orientation is wrong, the direction of down is wrong and a part of the gravitational acceleration will instead be believed to originate from the user moving the sensor. This error is double integrated to estimate the sensor position, resulting in a position error that grows cubically in time. This rapidly causes very large positioning errors. The gyro drift in a foot mounted IMU can be corrected if we can detect that the foot is on the ground. Then the foot is stationary and the gyro should be showing zero angular velocity. Instead it is showing the drift which can be estimated and then compensated for. This is known as ZUPT and reduces the positioning error to being linear in time [3].

Previously, stand still detection has been performed *ad hoc*, usually by comparing the signal to a threshold. In this work we put the stand still detection in a probabilistic framework using test statistics with known dis-

tributions and a Hidden Markov Model (HMM). The result is a probability of stand still at every time instant which can be used for ZUPT in a filtering framework.

Figure 1 shows an illustrative example of a map aided localization experiment without ZUPT. No stand still detections results in very uncertain step length estimates causing the position hypotheses to spread all along the corridor.

2 Related Work

Most solutions to the stand still detection problem use an averaged accelerometer or gyro value and compare it to a threshold [1, 2, 3, 5]. The threshold is chosen *ad hoc* and is normally quite low to minimize false positives. Another approach is the moving variance used in [4] where the variance computed over a sliding window is compared to a threshold.

Probabilistic zero velocity detection has previously been proposed in [6] who used a hypothesis test to determine if the foot was stationary or moving. The hypothesis test was performed using a test statistic based on a Generalized Likelihood Ratio Test (GLRT). The pdf of the acceleration and/or the angular velocity during the swing phase of the step, was approximated with an unnormalized uniform distribution. The pdf during stand still was based on the exponential of the norm of the acceleration and/or the angular velocity, which has an unknown distribution. The resulting test statistic was a moving average of the norm of the acceleration measurements and/or the angular velocity measurements. This was compared to a threshold to determine if the foot was to be rendered stationary. Since the test statistic has an unknown distribution the threshold was chosen *ad hoc*, making the framework similar to the ones in [1, 2, 3, 5].

The test statistics used in [6] are similar to the ones used in this work since we both evaluate three different ones where one is acceleration based, one is angular velocity based and one is based on a combination of acceleration and angular velocity. The acceleration based test statistic differs though in that we have chosen one which has a known distribution. This is also the case in the combined test statistic. Our framework to determine the mode probabilities also differs.

3 Stand Still Detection

The sensor is an Xsens MT motion sensor sampling at 100 Hz. The signals used are the accelerometers and the gyros. An example of a walking sequence is shown in Figure 2. The foot is stationary in the time instants around 550, 660, 770, 870 and 980. During these phases the norm of the accelerometer signals is the gravitation constant with noise. Simultaneously, the norm of the angular velocity signal is zero with some additive noise.




Figure 2: Example of accelerometer data (where x , y and z is blue, green and red) and gyro data (ω_x , ω_y and ω_z is blue, green and red, ω_i is angular rotation rate around axis i) during a walking sequence. The foot is stationary around time instants 550, 660, 770, 870 and 980.

3.1 Sensor Models

The signal model is

$$\mathbf{y}_t = \begin{bmatrix} \mathbf{y}_t^a(\theta) \\ \mathbf{y}_t^\omega(\theta) \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^a \\ \mathbf{v}_t^\omega \end{bmatrix} \quad (1)$$

where $\mathbf{y}_t^a$ and $\mathbf{y}_t^\omega$ denote acceleration vector and angular velocity vector, respectively. Further, θ denotes the model dependence of the phase of the human step sequence. Naturally, the model differs significantly between when the foot is at a stand still and when it is swinging.

The measurements have additive Gaussian noise distributions $\mathbf{v}_t^a \sim N(\mathbf{0}, \sigma_a^2)$ and $\mathbf{v}_t^\omega \sim N(\mathbf{0}, \sigma_\omega^2)$. The noise covariances are independent, resulting in $\sigma_\omega^2 = \sigma_\omega^2 \mathbf{I}$ and $\sigma_a^2 = \sigma_a^2 \mathbf{I}$, where $\mathbf{I}$ is the 3x3 identity matrix.

During stand still the sensor model can be described as

$$\begin{bmatrix} \mathbf{y}_t^a \\ \mathbf{y}_t^\omega \end{bmatrix} = \begin{bmatrix} g\mathbf{u} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^a \\ \mathbf{v}_t^\omega \end{bmatrix}, \quad (2)$$

where $\mathbf{u}$ is a gravitational direction vector and g is the gravitational constant 9.81. When the foot is moving the sensor model changes to

$$\begin{bmatrix} \mathbf{y}_t^a \\ \mathbf{y}_t^\omega \end{bmatrix} = \begin{bmatrix} g\mathbf{u} + \mathbf{a}_t \\ \boldsymbol{\omega}_t \end{bmatrix} + \begin{bmatrix} \mathbf{v}_t^a \\ \mathbf{v}_t^\omega \end{bmatrix} \quad (3)$$

where $\mathbf{a}_t$ and $\boldsymbol{\omega}_t$ have unknown distributions.

The problem is to safely distinguish between these two modes, stand still and swing. It is most important to minimize the stand still false positives, i.e. calling a stand still when the foot is in midair.

3.2 Test Statistic

To be able to differentiate between the two modes, test statistics with known distributions are computed. Three different ones are evaluated, one using only the accelerometer data, T^a , one using only the angular velocity data, T^ω , and one using a combination of both accelerometer and angular velocity data, T^c .

3.2.1 Acceleration Magnitude Detector

The accelerometer magnitude detector test statistic is computed as

$$T_t^a = \frac{\|\mathbf{y}_t^a\|^2}{\sigma_a^2} \quad (4)$$

where $T^a \sim \chi^2(3, \lambda)$ during stand still. It has a noncentral chi-square distribution since $\mathbf{y}_t^a$ has nonzero mean when the foot is stationary. Its noncentrality parameter $\lambda = g^2/\sigma_a^2$ and 3 is the number of degrees of freedom.

3.2.2 Angular Rate Magnitude Detector

The angular velocity test statistic is

$$T_t^\omega = \frac{\|\mathbf{y}_t^\omega\|^2}{\sigma_\omega^2} \quad (5)$$

where $T^\omega \sim \chi^2(3)$ during stand still since $\mathbf{y}^\omega$ has zero mean when the foot is stationary.

3.2.3 Combined Acceleration and Angular Rate Detector

The last test statistic combines acceleration and angular velocity to incorporate more information. It is calculated as

$$T_t^c = \frac{\|\mathbf{y}_t^a\|^2}{\sigma_a^2} + \frac{\|\mathbf{y}_t^\omega\|^2}{\sigma_\omega^2} \quad (6)$$

where $T^c \sim \chi^2(6, \lambda)$ during stand still. λ is the same as in (4) but the number of degrees of freedom has doubled to 6.

3.3 Test Statistic Appearance during Walking Sequence

A plot of the test statistics of the walking sequence in Figure 2 can be seen in Figure 3. The stand still events occurring around time instants 550, 660, 770, 870 and 980 are clearly visible. Figure 4 shows a zoom in of the test statistic with the mean of the stand still distribution marked with a dotted line. This reveals some of the problems with using only acceleration for stand still detection.

The test statistic T^a has a movement distribution that has a significant overlap of the stand still distribution, causing the test statistic to cross the mean of the stand still distribution during the stride. This is shown around time instants 530, 615, 630, 710, 750, 825 and 935. Simply calling a stand still when T^a is close to the mean of the stand still distribution will therefore cause a lot of false positives.




Figure 3: Test statistic from top to bottom; T_t^a , T_t^ω and T_t^c . The foot is stationary around time instants 550, 660, 770, 870 and 980.




Figure 4: Zoom in of the test statistics with the mean of the stand still distribution marked with a dotted line. The foot is stationary around time instants 550, 660, 770, 870 and 980.

T^ω has a distribution during movement that does not have a significant overlap of the stand still distribution, making T^ω a safer test statistic than T^a to use for stand still detection. Still, there are two occasions during the stride where the foot is quite stationary considering the angular velocity; one just after the foot has been lifted, in Figure 4 shown around time instants 600, 710, 815 and 920, and one just before set down shown at time instants 525, 635, 745 and 960. These can result in false positives.

The third test statistic T^c combines the strengths of T^a and T^ω . The bottom plot in Figure 4 shows that the foot does not appear stationary during the stride when you look at acceleration and angular velocity simultaneously. This results in a robust stand still detection.

4 Test Statistic Distribution Validation

The test statistics must be validated to ensure that the distribution of the test statistic under experimental stand still is close to the theoretical stand still distribution. We also estimate the distribution of the test statistic under experimental movements to illustrate the empirical probability density functions of stand still and movement that need to be separated.

4.1 Acceleration Magnitude Detector

The distributions of the acceleration magnitude test statistic T^a is shown in Figure 5. The theoretical and the empirical stand still distributions have similar mean but slightly different covariances. One of the reasons why the empirical density has a smaller covariance than the theoretical one, could be that we have been a bit too meticulous selecting the stand still data. Note also the significant overlap of the probability distributions of stand still and movement. That makes it difficult to safely identify stand stills by only looking at T^a .




Figure 5: Theoretical stand still distribution of T^a , empirical estimate of stand still distribution of T^a from experimental data and empirical estimate of movement distribution of T^a from experimental data.

4.2 Angular Rate Magnitude Detector

The distributions of T^ω is shown in Figure 6. Clearly, the theoretical stand still distribution is very similar to the empirical one, estimated by experimental data. Also note the large separation in magnitude of the empirical stand still and moving distributions. This enables more robust stand still detection than the distributions of T^a .




Figure 6: Theoretical stand still distribution of T^ω , empirical estimate of stand still distribution of T^ω from experimental data and empirical estimate of movement distribution of T^ω from experimental data.

4.3 Combined Acceleration and Angular Rate Detector

The combined test statistic naturally has distributions that look like combinations of the distributions of T^a and T^ω . The empirical stand still distribution has a similar mean but a slightly smaller covariance compared to the theoretical distribution. The empirical movement distribution does not overlap the stand still distributions as much as for T^a , enabling safer stand still detection.




Figure 7: Theoretical stand still distribution of T^c , empirical estimate of stand still distribution of T^c from experimental data and empirical estimate of movement distribution of T^c from experimental data.

5 Hidden Markov Model

To determine the probability of stand still, a Hidden Markov Model (HMM) is used. It determines the probability of each mode using the test statistic, the probability density functions of the modes and the mode transition probability matrix.

There are two modes; mode 1 when the foot is at a stand still and mode 2 when the foot is moving. The mode transition probability matrix states the probability of a mode switch which induces some dynamics into the probability estimation. A lower mode transition probability requires a measurement with a higher likelihood for the other mode to induce a switch.

The mode transition probability matrix is

$$\Pi = \begin{bmatrix} 0.95 & 0.05 \\ 0.05 & 0.95 \end{bmatrix} \quad (7)$$

which states that the probability of going from stand still to moving or vice versa, is 5%. During normal walking your right foot takes about one step per second which results in roughly 2 mode transitions every 100 measurements. The transition probabilities were chosen slightly higher to incorporate also faster movements.

The mode probabilities at time t are calculated using the recursion

$$\begin{aligned} \mu_t^i &= P(r_t = i | \mathbf{y}_t) \\ &\propto p(\mathbf{y}_t | r_t = i) P(r_t = i | \mathbf{y}_{t-1}) \\ &= p(T_t | r_t = i) \sum_{j=1}^{N_r} \Pi_{ji} \mu_{t-1}^j. \end{aligned} \quad (8)$$

Hence we have

$$\mu_t^i = \frac{p(T_t | r_t = i) \sum_{j=1}^{N_r} \Pi_{ji} \mu_{t-1}^j}{\sum_{l=1}^{N_r} p(T_t | r_t = l) \sum_{j=1}^{N_r} \Pi_{jl} \mu_{t-1}^j}. \quad (9)$$

The probability density function of movement used in the HMM is an approximation that is set to resemble the empirical movement density functions in Figures 5, 6 and 7. The HMM framework thus gives the probability of movement and stand still for each time instant. This framework can be extended to other modes like running and crawling, simply by extending the mode transition probability matrix by incorporating these new modes and estimating the probability densities for these movements too.

6 Experimental Results

The mode probabilities provided by the HMM of the data sequence in Figure 2 is shown in Figure 8. All three test statistics have been used to illustrate the difference in stand still detection performance.

The acceleration based test statistic T^a suffers from false positives around some of the troublesome time instants mentioned in Section 3.3; 615, 710, 750, 825 and




Figure 8: Mode probabilities for the different test statistics, evaluated on the data set in Figure 2. The foot is stationary around time instants 550, 660, 770, 870 and 980.

935. The framework does not call it a stand still after the first troublesome measurement, but after a couple of measurements close to the stand still mean the HMM assumes the foot is at rest.

The angular velocity based test statistic T^ω gives a distinct detection of every foot stance. The stationary moment is rather short but is often followed by a shorter second stationary moment. Figure 4 shows that this is because there is commonly a slight angular movement halfway through the deemed stationary part. This second stationary moment provides no new information and only the first detection is necessary to perform ZUPT. No false positives occur during the stride phase of the step.

The combined test statistic T^c provides very safe stand still detections. A long interval when the foot is at rest is deemed stationary and there are no false positives.

A second data set is constituted of a running phase followed by a walking phase, see Figure 9. The subject is running up until around time instant 1000. The foot is stationary around time instants 725, 810, 900, 1015 and 1130, the last two are during walking. The mode probabilities provided by the HMM of this sequence is shown in Figure 10. The same movement distribution was used during this whole experiment.

T^a does not provide any reliable stand still detections. The foot stances are detected, but a lot of false positives are also present. This is not surprising considering the accelerometer data in Figure 9. The gyro based T^ω does not result in any stand still detections at all during the running phase. This is a bit surprising since the gyro data looks pretty comprehensible and is probably because the IMU was fastened on the side of the foot where only very short periods of low angular




Figure 9: Example of accelerometer and gyro data during a sequence containing running, 650 – 1000, followed by walking, 1000 – 1150. The foot is stationary around time instants 725, 810, 900, 1015 and 1130. The data has the same color encoding as in Figure 2.

velocity are experienced.

The combined test statistic T^c still provides quite safe stand still detections. It picks up all the stand still sequences revealed by acceleration but manages to disregard the false ones using the angular velocity measurements. Here, the combined test statistic has shown to provide the most robust stand still detection.

Further walking experiments reveal the stand still detection performance shown in Table 1. All 174 true stationary phases were detected, but also some false positives. The acceleration based test statistic has a false positive between pretty much every step. Most false positives of the T^c statistic occur during sequences when small movements are performed like when a door is opened. During walking, T^ω gives very few false positives and is the safest stand still detector.

	T^a	T^ω	T^c
Stand stills detected	174	174	174
False positives	169	2	15

Table 1: True detected stationary phases and false detected stationary phases. 174 steps were taken.

6.1 Step Length Estimation

The foot mounted IMU has a coordinate system following the moving foot. In order to estimate the step length, the orientation of the foot in world coordinates is described by the unit quaternion $\mathbf{q}$. This relates the measured accelerations $\mathbf{y}_t^a$ and angular velocities $\mathbf{y}_t^\omega$ to movements and heading changes in the world coordinates. A filter with the states $\mathbf{p}$ and $\mathbf{v}$ for position and velocity in world coordinates and $\mathbf{q}$ can now be used




Figure 10: Mode probabilities for the different test statistics during the running followed by walking sequence. The foot is stationary around time instants 725, 810, 900, 1015 and 1130.

to estimate the length of each step. For a thorough description of the dynamical model, see [7].

A short experiment of 6 steps covering 5.0 meters was performed to evaluate the step length estimation performance. Stand still was detected using the gyro based test statistic and ZUPT was performed. Figure 11 shows the estimated movement in world coordinates. Total step length was estimated as 4.7 meters rendering a step length estimation error of 6%.




Figure 11: Movement estimation in world coordinates. 6 short steps were taken with a total length of 5.0 meters.

7 Conclusions and Future Work

Three test statistics with known distributions have been evaluated for stand still detection. The one based on accelerometer data only, has been shown to provide

plenty of false detections. This is natural since there is a significant overlap between the test statistic pdf during stand still and the pdf during movements. The gyro based has been shown to provide excellent stand still detection capabilities during walking while the one combining accelerometer and gyro data has been shown to provide good stand still detection during both walking and running. In conjunction with a Hidden Markov Model, the mode probabilities are readily calculated and can be used for zero velocity updates.

Future work includes extending the HMM framework to incorporate more modes and to merge the stand still detection with our localization framework. We will also look into whether the stationary phases detected using T^c are unnecessarily long for ZUPT. What we want to detect is gyro drift when the gyro is stationary and what we detect is when the combination of gyro and accelerometer is stationary. It is not necessarily the same thing. Further research is needed to decide when to perform the zero velocity update based on T^c .

References

- [1] S. Beauregard. Omnidirectional pedestrian navigation for first responders. In *Proc. of the 4th Workshop on Positioning, Navigation and Communication, WPNC07*, Hannover, Germany, 2007.
- [2] R. Feliz, E. Zalama, and J. G. Garcia-Bermejo. Pedestrian tracking using inertial sensors. *Journal of Physical Agents*, 3(1):35–43, 2009.
- [3] E. Foxlin. Pedestrian tracking with shoe-mounted inertial sensors. *IEEE Computer Graphics and Applications*, 25(6):38–46, 2005.
- [4] S. Godha, G. Lachapelle, and M. E. Cannon. Integrated GPS/INS system for pedestrian navigation in signal degraded environment. In *Proc. of ION GNSS*, 2006.
- [5] L. Ojeda and J. Borenstein. Non-GPS navigation for security personnel and first responders. *Journal of Navigation*, 60(3):391–407, 2007.
- [6] I. Skog. *Low-cost Navigation Systems*. PhD thesis, KTH, Stockholm, Sweden, 2009.
- [7] D. Törnqvist. *Estimation and Detection with Applications to Navigation*. PhD thesis, Linköping University, Linköping, Sweden, 2008.
- [8] O. Woodman and R. Harle. Rf-based initialisation for inertial pedestrian tracking. In *Proceedings of the 7th International Conference on Pervasive Computing*, 2009.