

Zero-Velocity Detection – A Bayesian Approach to Adaptive Thresholding

Johan Wahlström¹, Isaac Skog^{2*}, Fredrik Gustafsson^{3**}, Andrew Markham¹, and Niki Trigoni¹

¹Department of Computer Science, University of Oxford, Oxford OX1 2JD, UK

²S3 Research AB, 116 29 Stockholm, Sweden

³Department of Electrical Engineering, Linköping University, 581 83 Linköping, Sweden

* Senior Member, IEEE

** Fellow, IEEE

Abstract—A Bayesian zero-velocity detector for foot-mounted inertial navigation systems is presented. The detector extends existing zero-velocity detectors based on the likelihood-ratio test, and allows, possibly time-dependent, prior information about the two hypotheses – the sensors being stationary or in motion – to be incorporated into the test. It is also possible to incorporate information about the cost of a missed detection or a false alarm. Specifically, we consider an hypothesis prior based on the velocity estimates provided by the navigation system and an exponential model for how the cost of a missed detection increases with the time since the last zero-velocity update. Thereby, we obtain a detection threshold that adapts to the motion characteristics of the user. Thus, the proposed detection framework efficiently solves one of the key challenges in current zero-velocity-aided inertial navigation systems: the tuning of the zero-velocity detection threshold. A performance evaluation on data with normal and fast gait demonstrates that the proposed detection framework outperforms any detector that chooses two separate fixed thresholds for the two gait speeds.

Index Terms—Zero-velocity updates, foot-mounted inertial navigation, indoor localization, posterior odds ratio, adaptive thresholding.

I. INTRODUCTION

Zero-velocity-aided inertial navigation is one of the most promising technologies for indoor positioning in environments without pre-installed infrastructure [1]–[3]. The basic concept is illustrated in Fig. 1. A key component of this technology is the zero-velocity detector, which identifies when the sensors are stationary and zero-velocity updates can be applied. Hence, a variety of detectors have been proposed, many of which can be derived as generalized likelihood ratio tests [4].

One of the problems with existing zero-velocity detectors is that the optimal detection threshold varies significantly with the pedestrian’s walking speed and motion mode (walking, jogging, running, etc.), the placement of the sensor, the type of shoe, and the walking surface. If the threshold on the likelihood ratio is too large, the detector will not be able to detect stationary instances when the user is running. If the threshold is too small, the detector will produce false zero-velocity instances [5]. Consequently, several methods for designing adaptive thresholds have been proposed. The conducted studies have mainly focused on gait speed variations.

The most common approach to adaptive thresholding is to first use some heuristic or ad-hoc solution for estimating or classifying the speed or motion mode of the user. Based on the result, the detector selects a threshold value that has been optimized, using ground truth data, for that specific speed or motion class [6]–[10]. However, other methods for robust zero-velocity detection under varying gait conditions have also been explored. The authors in [11] used accelerometer measurements to detect the beginning and end of individual steps, and then applied different detectors and thresholds to different parts of the gait cycle; [12] made the detection using a long short-term memory (LSTM) neural network; [13] designed zero-velocity detectors that used a hidden Markov model (HMM) to represent the different stages of the gait cycle; [14] allowed the threshold to vary with the temporal variance of the accelerometer measurements; and [15] held the threshold fixed while dynamically




Fig. 1. Illustration of the concept of zero-velocity-aided inertial navigation. Using foot-mounted inertial sensors, the motion of the foot is tracked. To mitigate the error growth, zero-velocity updates are used to correct the navigation solution. One of the remaining challenges is how to reliably detect the zero-velocity instances during different usage and gait conditions.

adapting the window length of the samples used to compute the detection statistic.

Unfortunately, it is not clear how these methods relate to the existing theory on zero-velocity detection using likelihood ratio tests. Moreover, the proposed methods tend to require large data sets collected at different gait speeds and motion modes to calibrate the threshold values or other design parameters. Additionally, methods that attempt to infer speed or motion mode are always limited by the accuracy of the speed estimation or motion mode classification that precede the threshold selection.

This paper presents a zero-velocity detector based on the posterior odds ratio. The proposed detection framework demonstrates how the threshold, used in established detectors based on the likelihood ratio, can be factorized as a product of (i) the inverse prior odds ratio, quantifying the prior probability of a zero-velocity detection; and (ii) a loss factor, quantifying the cost of incorrect detections. The primary contributions are:

- A theoretical justification of adaptive zero-velocity detection within the established framework for zero-velocity detection based on the likelihood ratio test.
- An application-specific interpretation of the threshold used in traditional methods for zero-velocity detection.

- A discussion of how to model the hypothesis prior and the loss factor that specify the detection threshold.

The performance of the proposed detector is illustrated using a large data set collected at varying gait speeds¹.

II. BAYESIAN ZERO-VELOCITY DETECTION

Consider the problem of determining whether an inertial measurement unit (IMU) is stationary or not given the measurements $\mathbf{z}_n \triangleq \{\mathbf{y}_k\}_{k=n}^{n+N-1}$, collected over N sampling instances. Here, \mathbf{y}_k denotes the inertial measurements at sampling instance k . The problem can be formalized as the binary classification problem of choosing between the two hypotheses

$$\begin{aligned} \mathcal{H}_0 &: \text{IMU is moving} \\ \mathcal{H}_1 &: \text{IMU is stationary.} \end{aligned} \quad (1)$$

In what follows, a method for performing the classification by minimizing the conditional risk is described.

A. Minimum-Error-Rate Classification

The performance of a Bayesian detector is quantified using the conditional risk [16]

$$R(\mathcal{H}_i|\mathbf{z}_n) \triangleq \sum_{j=0}^1 \lambda_{ij} p(\mathcal{H}_j|\mathbf{z}_n). \quad (2)$$

The conditional risk $R(\mathcal{H}_i|\mathbf{z}_n)$ is the expected incurred loss when deciding on hypothesis \mathcal{H}_i given the data \mathbf{z}_n . Here, λ_{ij} represents the loss incurred as a result of deciding on hypothesis \mathcal{H}_i when the true hypothesis is \mathcal{H}_j . To minimize the conditional risk, we should decide on hypothesis \mathcal{H}_1 if and only if

$$\frac{p(\mathcal{H}_1|\mathbf{z}_n)}{p(\mathcal{H}_0|\mathbf{z}_n)} > \eta \quad (3)$$

where $\eta \triangleq (\lambda_{10} - \lambda_{00})/(\lambda_{01} - \lambda_{11})$. Throughout the paper, all denominators will be assumed to be nonzero.

B. Relation to Established Zero-Velocity Detectors

By using $p(\mathcal{H}_0) = 1 - p(\mathcal{H}_1)$ and the factorization

$$\frac{p(\mathcal{H}_1|\mathbf{z}_n)}{p(\mathcal{H}_0|\mathbf{z}_n)} = \underbrace{\frac{p(\mathbf{z}_n|\mathcal{H}_1)}{p(\mathbf{z}_n|\mathcal{H}_0)}}_{\text{Posterior odds ratio}} \cdot \underbrace{\frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)}}_{\text{Likelihood ratio}} \cdot \underbrace{\frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)}}_{\text{Prior odds ratio}}, \quad (4)$$

the detection rule in (3) can be written as

$$L(\mathbf{z}_n) \triangleq \frac{p(\mathbf{z}_n|\mathcal{H}_1)}{p(\mathbf{z}_n|\mathcal{H}_0)} > \gamma \quad (5)$$

where the detection threshold is

$$\begin{aligned} \gamma &\triangleq \frac{1 - p(\mathcal{H}_1)}{p(\mathcal{H}_1)} \cdot \eta \\ &= \frac{1 - p(\mathcal{H}_1)}{p(\mathcal{H}_1)} \cdot \frac{\lambda_{10} - \lambda_{00}}{\lambda_{01} - \lambda_{11}}. \end{aligned} \quad (6)$$

Thus, performing a zero-velocity detection using the posterior odds ratio in (3) is equivalent to performing a likelihood ratio test with a threshold dependent on the hypothesis prior $p(\mathcal{H}_1)$ and the loss factor η .

As shown in [4], many commonly applied zero-velocity detectors can, given different assumptions about the prior knowledge of the sensor signals, be derived as (generalized) likelihood ratio tests of the same form as (5). These include the acceleration-moving variance detector, the acceleration-magnitude detector, the angular rate energy detector, and the stance hypothesis optimal detection (SHOE) detector. Hence, the threshold in these detectors can be interpreted using (6). What is more, the relationship in (6) provides a theoretically sound way to design adaptive thresholds for these detectors. Next, we discuss how this can be done by letting the hypothesis prior $p(\mathcal{H}_1)$ and the loss factor η be time-dependent.

C. Modeling the Loss Factor

First, we set the cost of deciding on the correct hypothesis to zero, i.e., $\lambda_{00} = 0$ and $\lambda_{11} = 0$. The cost of a false alarm (incorrectly deciding that the IMU is stationary) is generally high due to the damaging effect of imposing an erroneous zero-velocity update on the navigation solution [5]. However, the problem of erroneous zero-velocity updates cannot easily be shown to be dependent on other factors, and hence, λ_{10} is assumed to be a constant. Moreover, since $\eta = \lambda_{10}/\lambda_{01}$, all time-dependence can without loss of generality be modeled in λ_{01} . The cost of a missed detection (incorrectly deciding that the IMU is moving) varies more distinctly with time. Immediately after a zero-velocity update, the uncertainty of the observable elements in the navigation state is low [17], and an additional zero-velocity update will only lead to a marginal improvement in estimation accuracy². However, the longer the time that has passed since the last zero-velocity update, the more important it is to break the error growth of the inertial navigation. Therefore, λ_{01} should increase with the time that has passed since the last detected zero-velocity instance. Further, since the likelihood ratio typically spans many orders of magnitude, a logarithmic or polynomial growth will seem slow and be hard to tune. Thus, in this paper, we will consider λ_{01} to have exponential growth. In summary, this means that the loss factor will be modeled to have an exponential decay according to

$$\eta = \alpha e^{-\theta \Delta t_k} \quad (7)$$

where α and θ are design parameters and Δt_k is the time since the last zero-velocity instance. If needed to avoid false alarms, it is possible to set a lower limit on η , so that $\eta = \max(\alpha e^{-\theta \Delta t_k}, \ell)$, where ℓ is some design parameter.

D. Modeling the Hypothesis Prior

The hypothesis prior will be dependent on what information is available in each specific scenario. In a zero-velocity-aided inertial navigation system implemented using a Bayesian filter, the system calculates a navigation solution in terms of a statistical distribution. This statistical distribution can be used to calculate a prior for the detector. Hence, consider the case when an extended Kalman filter is used, which at sampling instant k provides the velocity estimate $\hat{\mathbf{v}}_k$ and velocity error covariance \mathbf{S}_k . A measure of how close the system is to have zero velocity, weighted by the uncertainty of the velocity estimate, is then given by $\xi_k = \hat{\mathbf{v}}_k^\top \mathbf{S}_k^{-1} \hat{\mathbf{v}}_k$. Following

¹Reproducible research: The data and the code used in the experiments are available under an open-source licence at www.openshoe.org.

²In-fact, as described in [5], zero-velocity updates themselves induce systematic errors into the system, and excessive use of them may deteriorate the system performance.

the ideas of logistic regression, this measure can be mapped to the probability of the system being stationary via the logistic function. Thus, the prior is then set as

$$p(\mathcal{H}_1) = \frac{1}{1 + e^{\beta_1 \xi_k + \beta_2}} \quad (8)$$

where β_1 and β_2 are design parameters.

For ξ_k to be a reliable measure of how close the system is to have zero velocity, the covariance \mathbf{S}_k must reflect the true uncertainty of the velocity estimate $\hat{\mathbf{v}}_k$. If this cannot be guaranteed, the uninformative prior $p(\mathcal{H}_1) = 1/2$ may be a better choice.

E. Parameter Selection

Given the cost factor in (7) and the prior in (8), the logarithm of the detection threshold becomes

$$\log \gamma_k = c_1 + c_2 \Delta t_k + c_3 \xi_k \quad (9)$$

where $c_1 = \beta_2 + \log \alpha$, $c_2 = -\theta$, and $c_3 = \beta_1$. With an uninformative prior, (9) still holds but with $c_1 = \log \alpha$, $c_2 = -\theta$, and $c_3 = 0$. The parameters c_1 , c_2 , and c_3 may be estimated using ground truth data. However, this requires training data. Therefore, in what follows, a semi-heuristic parameter selection method is presented.

When the IMU is perfectly stationary the detector should with high probability decide on hypothesis \mathcal{H}_1 . This can be achieved by enforcing the condition $p(\log L(\mathbf{z}_n^*) < c_1) = \epsilon$, where ϵ is small (e.g., $\epsilon = 0.05$) and \mathbf{z}_n^* is the IMU output when the IMU is perfectly stationary. Here, we have used that ξ_k tends to be small at stationarity. For illustration, examine the first second in Fig. 2 where frequent zero-velocity instances are detected for a stationary IMU.

Similarly, it is reasonable to expect a zero-velocity update after completing a step. In mathematical terms, this gives us $p(\log L(\mathbf{z}_n^*) < c_1 + c_2 \Delta \tau) = \epsilon$, where $\Delta \tau$ is the approximate time length of a step and \mathbf{z}_n^* is the output from the IMU during the midstance of normal gait. In this way, most steps during normal walking will result in detected zero-velocity instances. Fig. 2 illustrates how the threshold and likelihood ratio intersect at the end of each step.

With an uninformative prior, one may select the parameters c_1 and c_2 by using the two conditions described above and set $c_3 = 0$. However, with the prior in (8) we need a third condition to perform the parameter selection. To this end, note that the detector should decide on hypothesis \mathcal{H}_0 during the swing phase. Therefore, we impose the condition $p(\log L(\mathbf{z}_n^*) < c_1 + c_2 \Delta \tau / 2 + c_3 \xi_k^*) = 1 - \epsilon$ where \mathbf{z}_n^* and ξ_k^* are typical values during the swing phase of \mathbf{z}_n and ξ_k , respectively. The effect of using the prior in (8) can be seen by comparing the dashed black line and the solid red line in Fig. 2.

Finally, note that although the proposed detection framework includes design parameters, it will not be as sensitive to parameter tuning as the conventional detector using a fixed threshold. Firstly, as outlined above, there are thumb rules for tuning the Bayesian detector. Secondly, regardless of how the Bayesian detector is tuned, the risk of missing zero-velocity instances when the user is running (which is the consequence of choosing a too large fixed threshold) is mitigated by the fact that the threshold diminishes with the time since the last detected zero-velocity instance.

III. EXPERIMENTS

The performance of the proposed detector was benchmarked against a detector with a fixed threshold using data from [18]. The data set

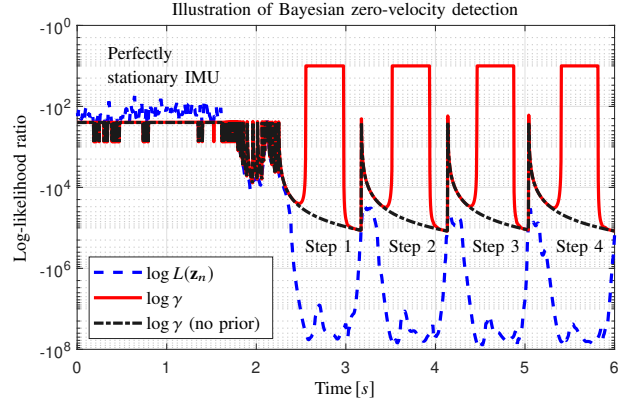


Fig. 2. Illustration of Bayesian zero-velocity detector using the likelihood ratio of the SHOE detector (see [4]) and an adaptive threshold.

was collected with a MicroStrain 3DM-GX2 IMU and consists of 40 data recordings; half with normal gait (5 [km/h]) and the other half with fast gait (7 [km/h]). In each recording, the user walked one lap along a closed-loop trajectory with an approximate length of 84 [m]. The zero-velocity detection was performed using the likelihood ratio of the SHOE detector (see [4]) with a window length of 20 [ms] (5 samples). The parameters c_1 , c_2 , and c_3 were quickly tuned based on the procedure outlined in Section II-E, and remained fixed on all data. The two hypothesis priors suggested in Section II-D resulted in the same navigation performance. However, for gait with longer periods between the zero velocity events, the informative prior suggested in Section II-D may prevent false zero-velocity detections otherwise caused by too aggressive growth of the cost of a missed detection.

Fig. 3 displays root-mean-square errors (RMSEs) computed from the final position estimates in each trajectory for both the Bayesian detector and the corresponding detector with a fixed threshold. The adaptive threshold is seen to provide a lower RMSE than the best fixed threshold. This holds true not only when performing the evaluation on all data, but also individually on the data with normal gait and the data with fast gait. Thus, these results indicate that the Bayesian detector proposed in this paper is preferable to previously investigated methods for adaptive thresholding that choose a separate threshold for each speed or motion mode, see e.g., [6]–[10]. To illustrate that the difference in performance becomes rather substantial over longer time periods, we also computed the position error resulting from concatenating all available data recordings (corresponding to walking a distance of about 3.4 [km]). In this case, the adaptive threshold and the best fixed threshold give position errors with magnitudes of 8.16 [m] and 20.35 [m], respectively.

The performance improvement may be explained in several ways. On the one hand, the adaptive threshold prevents the navigation system from running too long without a zero-velocity update, and thereby prevents the estimation errors from growing too large before the next zero-velocity update. On the other hand, the adaptive threshold prevents the navigation system from using an excessive number of zero-velocity updates, and thereby reduces the negative impact that modeling errors in the pseudo velocity measurements have on the navigation solution.

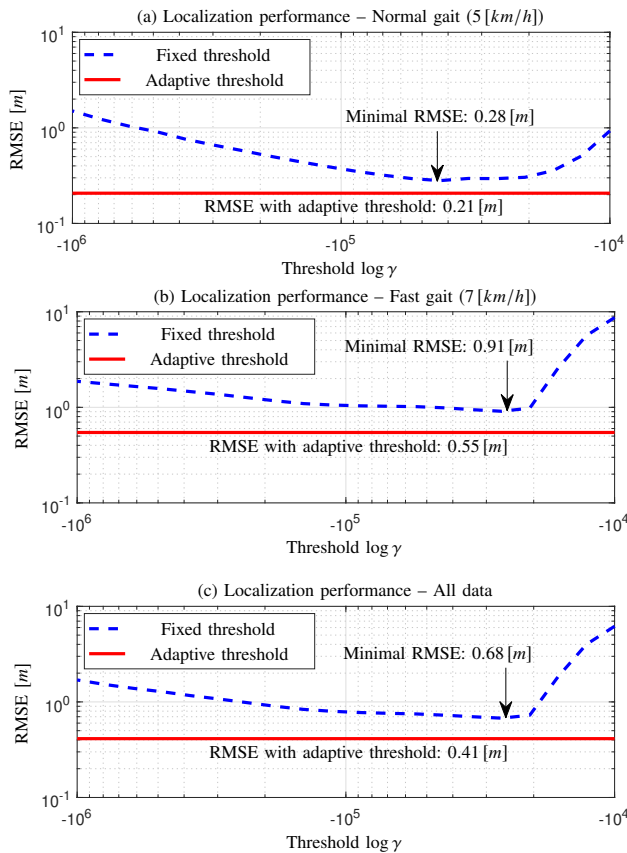


Fig. 3. Position error of zero-velocity-aided inertial navigation using a fixed and an adaptive threshold.

IV. CONCLUSION AND DISCUSSION

This paper has developed a Bayesian zero-velocity detector using the posterior odds ratio. Experiments using data collected at varying gait speeds indicated that the proposed detector provides a significantly lower positioning error than state-of-the-art detectors where an adaptive threshold is selected based on the motion mode of the user. Further, the proposed method does not rely on learning the relationship between the user motion and the optimal threshold, and therefore avoids the time-consuming data collection that has been necessary in previous methods for adaptive thresholding.

One possible direction for future work is to evaluate the Bayesian detector for a broader set of motion classes and gait conditions. Although our results indicate that the Bayesian detector alleviates the need for adaptive parameter tuning, the Bayesian detector could, in the same way as the conventional detector with a fixed threshold, still be used together with motion classifiers that enable adaptive parameter tuning.

ACKNOWLEDGMENT

This research has been financially supported by the National Institute of Standards and Technology (NIST) via the grant *Pervasive, Accurate, and Reliable Location-based Services for Emergency Responders* (Federal Grant: 70NANB17H185).

REFERENCES

- [1] A. Norrdine, Z. Kasmi, and J. Blankenbach, "Step detection for ZUPT-aided inertial pedestrian navigation system using foot-mounted permanent magnet," *IEEE Sensors J.*, vol. 16, no. 17, pp. 6766–6773, Sep. 2016.
- [2] M. N. Muhammad, Z. Salcic, and K. I. Wang, "Indoor pedestrian tracking using consumer-grade inertial sensors with PZTD heading correction," *IEEE Sensors J.*, vol. 18, no. 12, pp. 5164–5172, Jun. 2018.
- [3] J. Wahlström, P. Porto Buarque de Gusmão, A. Markham, and N. Trigoni, "Map-aided navigation for emergency searches," in *Proc. IEEE Int. Conf. Distrib. Comput. in Sensor Syst.*, Santorini Island, Greece, May 2019.
- [4] I. Skog, P. Händel, J. O. Nilsson, and J. Rantakokko, "Zero-velocity detection — An algorithm evaluation," *IEEE Trans. Biomed. Eng.*, vol. 57, no. 11, pp. 2657–2666, Nov. 2010.
- [5] J.-O. Nilsson, I. Skog, and P. Händel, "A note on the limitations of ZUPTs and the implications on sensor error modeling," in *Proc. IEEE Int. Conf. Indoor Positioning and Indoor Navigation*, Sydney, Australia, Nov. 2012.
- [6] X. Tian, J. Chen, Y. Han, J. Shang, and N. Li, "A novel zero velocity interval detection algorithm for self-contained pedestrian navigation system with inertial sensors," *Sensors*, vol. 16, no. 10, p. 1578, Jun. 2016.
- [7] S. Y. Park, H. Ju, and C. G. Park, "Stance phase detection of multiple actions for military drill using foot-mounted IMU," in *Proc. IEEE Int. Conf. Indoor Positioning and Indoor Navigation*, Alcalá de Henares, Spain, Oct. 2016.
- [8] R. Zhang, H. Yang, F. Höflinger, and L. M. Reindl, "Adaptive zero velocity update based on velocity classification for pedestrian tracking," *IEEE Sensors J.*, vol. 17, no. 7, pp. 2137–2145, Apr. 2017.
- [9] M. Ma, Q. Song, Y. Li, and Z. Zhou, "A zero velocity intervals detection algorithm based on sensor fusion for indoor pedestrian navigation," in *Proc. IEEE Int. Conf. Inf. Technol., Netw., Electron. Autom. Control*, Chengdu, China, Dec. 2017, pp. 418–423.
- [10] B. Wagstaff, V. Peretroukhin, and J. Kelly, "Improving foot-mounted inertial navigation through real-time motion classification," in *Proc. IEEE Int. Conf. Indoor Positioning and Indoor Navigation*, Sapporo, Japan, Sep. 2017.
- [11] Y. Li and J. J. Wang, "A robust pedestrian navigation algorithm with low cost imu," in *Proc. IEEE Int. Conf. Indoor Positioning and Indoor Navigation*, Sydney, Australia, Nov. 2012.
- [12] B. Wagstaff and J. Kelly, "LSTM-based zero-velocity detection for robust inertial navigation," in *Proc. IEEE Int. Conf. Indoor Positioning and Indoor Navigation*, Nantes, France, Sep. 2018.
- [13] W. Sun, W. Ding, H. Yan, and S. Duan, "Zero velocity interval detection based on a continuous hidden markov model in micro inertial pedestrian navigation," *Meas. Sci. Technol.*, vol. 29, no. 6, Apr. 2018.
- [14] H.-F. Liu, W. Ren, T. Zhang, J. Gong, J. m. Liang, B. Liu, J. w. Shi, and Z. Huang, "An adaptive selection algorithm of threshold value in zero velocity updating for personal navigation system," in *Proc. 33rd Chinese Control Conf.*, Jul. 2014, pp. 1035–1038.
- [15] Q. Wang, Z. Guo, Z. Sun, X. Cui, and K. Liu, "Research on the forward and reverse calculation based on the adaptive zero-velocity interval adjustment for the foot-mounted inertial pedestrian-positioning system," *Sensors*, vol. 18, no. 5, p. 1642, May 2018.
- [16] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*. Wiley, 2000.
- [17] J.-O. Nilsson, D. Zachariah, I. Skog, and P. Händel, "Cooperative localization by dual foot-mounted inertial sensors and inter-agent ranging," *EURASIP J. Advances in Signal Process.*, no. 1, Oct. 2013.
- [18] I. Skog, J. O. Nilsson, and P. Händel, "Evaluation of zero-velocity detectors for foot-mounted inertial navigation systems," in *Proc. IEEE Int. Conf. Indoor Positioning and Indoor Navigation*, Zurich, Switzerland, Oct. 2010.