Probabilistic Tracking of Pedestrian Movements via In-Floor Force Sensing

Rishi Rajalingham, Yon Visell, Jeremy R. Cooperstock
Centre for Intelligent Machines and CIRMMT, McGill University
rishi@cim.mcgill.ca, yon@cim.mcgill.ca, jer@cim.mcgill.ca

Abstract

This article presents a probabilistic approach to the tracking and estimation of the lower body posture of users moving on foot over an instrumented floor surface. The latter consists of an array of low-cost force platforms providing intermittent foot-floor contact data with limited spatial resolution. We use this data to track body posture in 3D space using Bayesian filters with a switching state-space model. Potential applications of this work to person tracking and human-computer interaction are described.

1. Introduction

3D human posture tracking is a classic challenge in computer vision and pattern recognition. Computer vision techniques have been the most widely used for this purpose [7, 8], although several researchers have investigated human tracking via in-floor sensor arrays [10, 11, 1]. Common challenges of person tracking in these domains include the loss of 3D pose information through observation via the sensor array, and the underlying complex dynamics of human motion. Missing information plays an important role in both settings. Losses due to occlusion in video-based tracking are somewhat analogous to the loss of observations while feet are out of contact with the ground during tracking via a floor-based array. Overcoming such missing information has, in part, motivated the approach presented here.

Although motion capture and video can provide high-resolution 3D position information for human tracking, they are not always available and are prone to visual occlusion. Furthermore, state of the art methods for inferring human contact interactions from video provide inaccurate information about interaction forces between body and ground [3]. Such forces are highly characteristic of the individuals and activities generating them (as evidenced in the references given below).

Bayesian filtering provides a unifying view of diverse probabilistic tracking methods. It has been extensively applied to problems in object, person, or context tracking [4, 17]. The effectiveness of such methods stems from their ability to integrate information acquired over time in ways that respect the structure and dynamics of the object or individual being tracked. Nonparametric Bayesian filtering techniques, like the particle filter based model used here, make it possible to perform tracking without making unnecessary assumptions about the form taken by those distributions or the dynamics governing them, by maintaining many hypotheses in parallel.

The system presented here consists of a posture tracking system based on a distributed, sparse in-floor sensing array. Prior literature on the analysis of foot-floor contact forces has addressed applications to immersive interactive media [12], pedestrian identification [11, 16], gait and dance analysis [13], among others [21, 15, 9, 6]. Here, similar data is used for 3D kinematic tracking of users' lower bodies, a task that has received less attention in prior research on floor-sensing arrays. Murakita et al. utilized a Markov Chain Monte Carlo method to track pedestrians via an array of binary pressure sensors [10], albeit with much larger errors (typically 0.6 m) than what we achieve here (Sec. 4). Yin and Pai tracked whole-body movements via a high-resolution (and costly) floor sensing array, based on the similarity of force patterns to those recorded in a database of movements [21]. However, their system was limited to tracking relative to a predefined set of static poses.

Our intended applications involve both person tracking and interaction with floor-based touch displays via the feet [20]. In both cases, we are interested in tracking the 3D kinematics of the lower body of moving persons, with an emphasis on the locations of their feet. Although existing force sensing arrays can provide accurate information about foot-floor contact forces, the achievable resolution is often limited by cost constraints. In addition, mappings from patterns of foot-floor forces to body posture are complex and one-to-many. Effective use of prior knowledge about body structure, movement, and walking mechanics is required in order to track posture accurately.

2. System Configuration

The sensing floor (Fig. 1) consists of a 6×6 array of rigid tiles, 30 cm on each side. Each tile is instrumented with four resistive force sensors (Interlink model 402), which are located at the corners. Thus, the nominal linear sensing resolution of this array is 15 cm. In addition, each tile includes a wide bandwidth vibrotactile actuator [18], which although unused here, does modestly influence the sensor measurements, due to its weight [20]. Analog data from the force sensors are conditioned, amplified, and digitized via 32-channel, 16-bit data acquisition boards. Each sensor is sampled at a rate of up to 1 kHz transmitted over a low-latency Ethernet link. An array of six small-form-factor

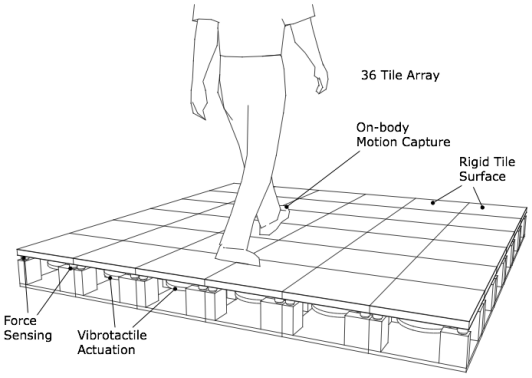

Figure 1. Illustration of the distributed floor interface, with components labelled.

computers is used for force data processing.

For applications that do not require kinematic tracking, we infer foot-floor contact loci using intrinsic contact sensing techniques [20], attaining an effective resolution of about 2 cm. However, such methods are incapable of tracking the location of a foot once it leaves the floor surface, and cannot resolve situations in which the feet overlap onto a single tile. Moreover, inference of lower body pose from in-floor force sensor data is challenging due to the loss of information inherent in this complex mapping. Such limitations have, in part, motivated the present work. In particular, we adopt the framework of Bayesian filtering in order to maintain continuity of lower body position estimates in dynamic settings, such as walking, where foot-ground contact is regularly interrupted.

Within this framework, our task remains challenging due to the high dimensional mapping from an articulated pedestrian pose to the observed force sensor values. In particular, this map is discontinuous for a sensing infrastructure consisting of independent, rigid tiles, as poses that are similar in nature, i.e. with footprints of similar 2D position and orientation, will lead to drastically different force observations when tile boundaries are crossed.

3. Tracking Problem and Algorithm

Consider a dynamic system with states x_t and observations z_t, both indexed by time. A Bayesian filter probabilistically estimate at time t the state x_t by sequentially updating a belief distribution $\mathrm{Bel}(x_t)$ over the state space, defined by $\mathrm{Bel}(x_t) = p(x_t|z_t, z_{t-1}, z_{t-2}, ...) = p(x_t|z_{1:t})$. Assuming the states comprise a Markov process, the Belief distribution at each subsequent time step can be obtained, using Bayes' Theorem, in terms of the belief state at the prior time step $t-1$, the assumed motion model $p(x_t|x_{t-1})$, and the likelihood $p(z_t|x_t)$ of the newly acquired observation z_t given the state x_t:

$$\mathrm{Bel}(x_t) \propto p(z_t|x_t) \int p(x_t|x_{t-1})\mathrm{Bel}(x_{t-1})\, dx_{t-1} \quad (1)$$

Bayesian filters can be distinguished, in part, by the form of the Belief distribution $\mathrm{Bel}(x_t)$ that is assumed. Here, we adopt a Monte-Carlo approach, in which $\mathrm{Bel}(x_t)$ is represented by a set of weighted samples, or particles, given by $S_t = \{(x_t^i, w_t^i), i = 1, 2, \ldots N_s\}$. Our method uses the Sampling-Importance-Resampling (SIR) algorithm [5], described in Alg. 1, and schematically illustrated in Fig 2.

For each new observation, SIR re-weighs the set of particles depending on their likelihood, evolves them in time using the assumed dynamic model, and updates the particle set S_t by resampling based on these weights.

Specifically, the likelihood function $L(x_t) = p(z_t|x_t)$ is computed in two steps as follows.

1. Using the observation model $\mathbf{H}$, defined in Section 3.1, that maps states into the observation space, generate expected observations $z_t^* = \mathbf{H}(x_t)$.

2. Using the similarity measure $S(z, z') = p(z|z')$, defined in Section 3.2, compute the likelihood as $L(x_t) \equiv S(z_t^*, z_t)$.

As seen in Eq.2, each particle is weighted according to its likelihood while also considering a prior distribution $p(x_t)$. As described in Section 3.2, the prior distribution $p(x_t)$ is useful for encoding constraints on states x_t, thus leading to an efficient search of the state space.

The attributed weight is then used to resample the particle set (see Eq. 3). This step distinguishes SIR from SIS (Sampling Importance Sampling), and is meant to eliminate degeneracy of particles; by concentrating on high weighted particles to create a new uniformly distributed particle set, the extreme case where most particles have negligible weight is avoided. The opposite extreme, consisting of a single strong hypothesis, can be avoided by the inclusion of a roughening process. Roughening consists of the addition of random, zero-mean, normally distributed noise to all particles after resampling (Eq. 5), to allow for a thorough search of the solution space. Additionally, particles

are propagated forward in time based on the assumed dynamic model (see Eq.4), which implicitly defines the motion probability $p(x_t|x_{t-1})$. Our dynamic model consists of two processes, accounting for the continuous and discrete aspects of the motion of interest. The dynamic model and roughening process are defined in detail in Section 3.3.

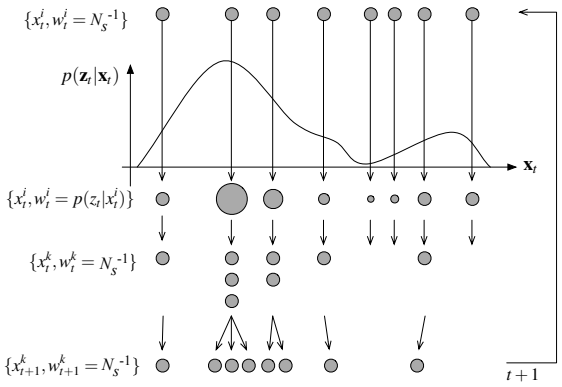

Figure 2. The SIR particle filter algorithm.

Algorithm 1 Sampling Importance Resampling

Initialize particles randomly: $S_0 \sim N(\mu_S, \sigma_S)$
while $t > 0$ **do**
Observe z_t.
for $i = 1$ to N_s **do**

$$\text{Likelihood:} \quad p(z_t|x_t^i) = S(z_t, H(x_t))$$
$$\text{Weight:} \quad w_t^i = p(x_t^i) \times p(z_t|x_t^i) \qquad (2)$$

end for
for $i = 1$ to N_s **do**

$$\text{Normalize:} \quad W = \textstyle\sum_i w_t^i, w_t^i \leftarrow W^{-1} w_t^i$$
$$\text{Resample:} \quad x_t^i \sim p(x_t^{i*}|w_t^i), w_{t+1}^i = N_s^{-1} \qquad (3)$$
$$\text{Draw:} \quad x_{t+1}^i \sim p(x_{t+1}|x_t^i) \qquad (4)$$
$$\text{Roughen:} \quad x_{t+1}^i \leftarrow x_{t+1}^i + \eta_{t+1}, \eta \sim N(0, \sigma_x) \qquad (5)$$

end for
$S_{t+1} = \{(x_{t+1}^i, w_{t+1}^i)\}_{i=1}^{N_s}$
end while

In our system for tracking via in-floor force measurements, the relevant variables consist of:

- Observations $\mathbf{z}_t$, consisting of a 12×12 array of force values, f_i.
- States, $\mathbf{x}_t$, describing kinematic lower-body poses, are 19-dimensional vectors: $\mathbf{x}_t$ = ($\phi_{l,t}$, $\dot{\phi}_{l,t}$, $\dot{\phi}_{l,t-1}$, $\phi_{r,t}$, $\dot{\phi}_{r,t}$, $\dot{\phi}_{r,t-1}$, β). They include planar midpoint coordinates $\mathbf{u}$ and orientations θ for each foot, where $\phi_l = (\mathbf{u}_l, \theta_l)$ and likewise for ϕ_r, along with first time derivatives. The state $\mathbf{x}_t$ also includes a binary-valued vector $\beta = (\beta_l, \beta_r)$, implemented as a quaternary variable, which indicates the foot-floor contact condition ($\beta_i = 1$ if there is contact) for the left and right feet.

Figure 3 illustrates this state description within a skeletal model. The algorithm definition is completed by specifying the observation model $\mathbf{z}_t = \mathbf{H}(\mathbf{x}_t)$, the likelihood model $p(\mathbf{z}_t|\mathbf{x}_t)$, and the motion model, $p(\mathbf{x}_t|\mathbf{x}_{t-1})$.

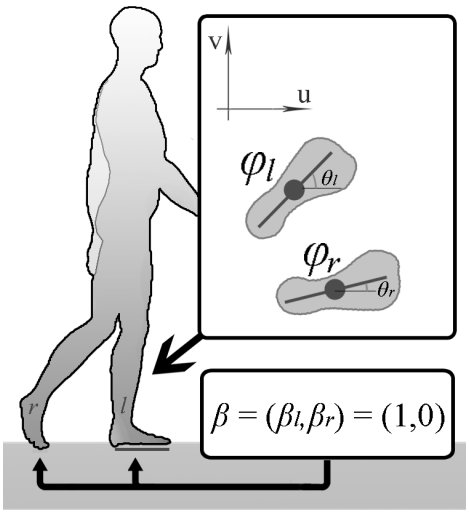

Figure 3. State description of lower body poses: $\mathbf{x}_t$ defines feet and foot-floor contact.

3.1 Observation model

We model expected observations $\mathbf{H}(\mathbf{x}_t)$ for a state $\mathbf{x}_t$ by simulating the mechanical forces associated with a pose. Ignoring shear forces, foot-tile contact results in a normal pressure distribution $p(\mathbf{u})$, where $\mathbf{u} = (u, v)$ are 2D coordinates on the floor. The pressure distribution on a tile is conveniently summarized by a contact centroid $p_c = (\mathbf{u}_c, F)$, where $F = \int d\mathbf{u}\, p(\mathbf{u})$ is the net normal force and $\mathbf{u}_c$ is the pressure centroid. A normal force with magnitude F, applied at $\mathbf{u}_c$, would give rise to the same force measurements f_i as $p(\mathbf{u})$ [20, 2] (Fig 4). Our observation model associates a pose $\mathbf{x}_t$ to a set of contact centroids, $p_{c,j}, j = 1, 2, \ldots, N_c$. One centroid is placed on each tile that a foot pose is determined to be in contact with ($\beta_i > 0$), and the total weight F of the user is partitioned among these centroids. The placement of a contact centroid is determined by the average position of foot to tile contact. Expected sensor readings f_i are obtained from the static equilibrium equations for each tile. Each pressure centroid p_c yields a contribution $f_{i,c} = d_i^{-1}(\sum_{j=0}^{4} d_j^{-1})^{-1}$, where $d_j = |\mathbf{u}_c - \mathbf{u}_j|$, and $\mathbf{u}_j$ is the sensor location. In this way, we predict observed force values $\mathbf{z}_t^* = \{f_{i,t}\}$ for a state $\mathbf{x}_t$.

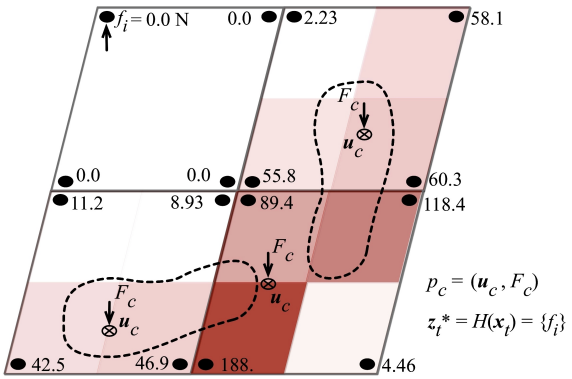

Figure 4. The observation model $H(\mathbf{x}_t)$ maps a pose $\mathbf{x}_t$ (illustrated here as two feet) to a set of contact centroids ($\mathbf{u}_c$). These are then converted to force sensor readings $\mathbf{z}_t = \{f_i\}$. Expected force observations are illustrated as grids, with each quadrant intensity proportional to corresponding force sensor value (given in Newtons).

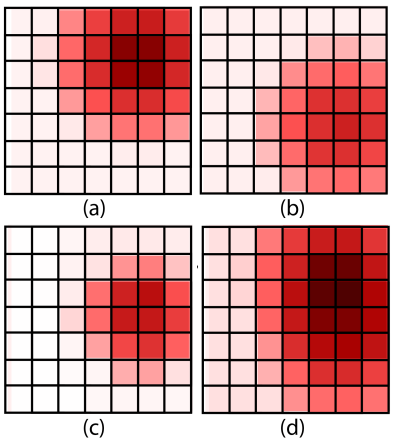

Figure 5. 2D pressure distributions for (a) $\mathbf{z}_t$, (b) $\mathbf{z}_t^*$, (c) $\mathbf{z}_t \cap \mathbf{z}_t^*$, (d) $\mathbf{z}_t \cup \mathbf{z}_t^*$. Pressure distributions are illustrated as grids, with each quadrant intensity proportional to corresponding force sensor value. The intersection ($\cap$) and union ($\cup$) of force observations are used in computing similarity $S(\mathbf{z}_t, \mathbf{z}_t^*)$.

3.2 Likelihood model

The likelihood function $L(\mathbf{x}_t) = p(\mathbf{z}_t|\mathbf{x}_t)$ describes the probability of observing force values $\mathbf{z}_t$ given the state $\mathbf{x}_t$. As described above, we use the observation model to generate expected observations, $\mathbf{z}_t^* = \mathbf{H}(\mathbf{x}_t)$. The likelihood function is then defined in terms of a normalized similarity measure $p(\mathbf{z}_t|\mathbf{x}_t) = S(\mathbf{z}_t^*, \mathbf{z}_t)$ between true force pattern observations $\mathbf{z}_t$ and expected force observations $\mathbf{z}_t^*$.

Similarity measure S The similarity measure $S(\mathbf{z}, \mathbf{z}') = p(\mathbf{z}|\mathbf{z}')$ models the probability of observing $\mathbf{z}$ if the true observation is $\mathbf{z}'$. Conventional similarity measures make use of metrics such as Euclidean or Mahalanobis distance functions. However, since $\mathbf{H}$ is a high-dimensional discontinuous map from states $\mathbf{x}_t$ to observations $\mathbf{z}_t$, these measures cannot properly gauge similarity between observation vectors $\mathbf{z}_t$, $\mathbf{z}_t^* = \mathbf{H}(\mathbf{x}_t)$. Moreover, the force observations are comparatively sparse, with most values being zero. As an alternative, we compute pair-wise similarity between such patterns, based on a measure of their area of overlap. Specifically, we employ a similarity measure that has proved useful in tracking via binary image masks [14], computing $S(\mathbf{z}^*, \mathbf{z})$ as the relative area of overlap between the true and expected 2D pressure distributions,

$$S(\mathbf{z}_t^*, \mathbf{z}_t) = \frac{\cap(\mathbf{z}_t^*, \mathbf{z}_t)}{\cup(\mathbf{z}_t^*, \mathbf{z}_t)} = \frac{1}{N_z} \sum_{i=1}^{N_z} \frac{\min(\mathbf{z}_t(i), \mathbf{z}_t^*(i))}{\max(\mathbf{z}_t(i), \mathbf{z}_t^*(i))}$$

Fig. 5 illustrates conceptual examples of observations $\mathbf{z}_t$, $\mathbf{z}_t^*$ as 2D pressure distributions, as well as their intersection $\cap(\mathbf{z}_t^*, \mathbf{z}_t)$ and union $\cup(\mathbf{z}_t^*, \mathbf{z}_t)$.

We note that the average overlap of these resulting pressure distributions is an effective metric for capturing similarities between force observations as a probability:

$$\begin{aligned} S(\mathbf{z}, \mathbf{z}) &= 1 \\ \mathbf{z}_1 \neq \mathbf{z}_2 &\Rightarrow 0 < S(\mathbf{z}_1, \mathbf{z}_2) < 1 \\ S(\mathbf{z}_1, \mathbf{z}_2) &= S(\mathbf{z}_2, \mathbf{z}_1) \end{aligned} \quad (6)$$

Postural constraints The likelihood model is modified to encode human postural constraints, via a prior distribution $p(\mathbf{x}_t)$. The latter is defined to consist of a set of independent postural priors for human walking, in the form of univariate Gaussian distributions $N(\mu_\varsigma, \sigma_\varsigma)$, $N(\mu_\Theta, \sigma_\Theta)$ over stance width $\varsigma = \|\mathbf{u}_l - \mathbf{u}_r\|$ and relative orientation $\Theta = |\theta_l - \theta_r|$ respectively. The postural prior is introduced as the product $p(\mathbf{x}_t) = N(\varsigma; \mu_\varsigma, \sigma_\varsigma) N(\Theta; \mu_\Theta, \sigma_\Theta)$. The prior distribution $p(\mathbf{x}_t)$ is applied when computing the particle weights from the likelihood L (see Eq. 2).

3.3 Dynamics Model

We model the movements of the lower body of a walker via the motion probability $p(\mathbf{x}_t|\mathbf{x}_{t-1})$. The state x_t consists of continuous configuration variables ϕ_i and discrete contact variables β_i. We therefore approximate foot motion in a hybrid (stochastic) framework, incorporating continuous, linear movements of the limb coupled to discrete state transitions reflecting changes of the foot-floor contact conditions.

Continuous, linear dynamics The state representing the left or right foot is denoted respectively by a vector $\phi_i = (\mathbf{u}_i, \theta_i)$ giving the midpoint and orientation of either the left ($i = l$) or right ($i = r$) foot at time t (Fig. 3. The dynamics used in our algorithm can be described by the following linear, discrete time system:

$$\phi_{i,t} = \phi_{i,t-1} + \beta_i \Gamma_i (\dot{\phi}_{t-1} dt + \eta_t^0) \quad (7)$$

$$\dot{\phi}_{i,t} = \beta_i [\alpha \dot{\phi}_{i,t-1} + (1-\alpha)(\dot{\phi}_{i,t-2} + \eta_{t-1}^1)] \quad (8)$$

The parameter β_i is the binary contact variable for foot i. Thus, ϕ_i is constant when there is contact ($\beta_i = 0$) and otherwise drifts, with position and velocity driven by additive Gaussian noise processes η_t^0 or η_t^1, where $\eta \sim N(0, \Sigma)$. For efficiency, we parametrize drift via a single noise process, defining $\eta \equiv \eta_t^0 = \eta_t^1 dt$ for all t. The noise covariance Σ is a 3×3 diagonal matrix with diagonal entries $\sigma_u, \sigma_v, \sigma_\theta$. To mimic walking, velocity drift in the direction that the foot is oriented is assumed to be larger. This non-isotropic drift is implemented through the factor Γ_i, a diagonal matrix with entries $(\cos\theta_i(\gamma + \sqrt{1-\gamma^2}), \sin\theta_i(\gamma - \sqrt{1-\gamma^2}), 1)$, where $0 < \gamma < 1$ is a dimensionless scalar defining the longitudinal bias. α is a dimensionless scalar defining the velocity noise mixing rate. It approximates the dynamics of a free foot during walking by means of a saturating linear drift velocity.

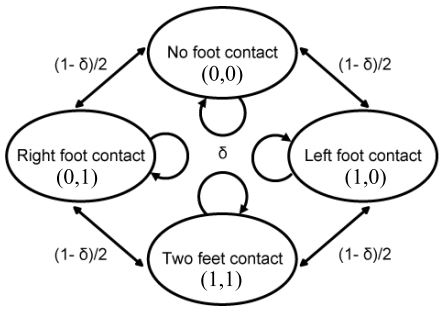

Figure 6. Stochastic state transition diagram approximating stepping motion.

Roughening As presented in (Alg. 1), noise $\eta \sim N(0, \sigma_x)$ is added to the continuous state components ϕ_i at each SIR step in order to avoid particle degeneracy.

Discrete state transition model The dynamic model includes discrete transitions from the foot-floor contact states $\beta = (\beta_l, \beta_r)$, where $\beta = 0$ or 1, via the stochastic process, shown in Figure 6. Despite its simplicity, this model is effective in approximating discrete stepping motion. δ is an empirically determined probability of no change in contact state β. All remaining transitions are symmetric, with transition probability $\frac{1-\delta}{2}$.

Foot	Error (m)	Windowed error (m)
Right (no contact)	0.1901	0.1245
Left (no contact)	0.1649	0.0336
Average (no contact)	**0.1775**	**0.0791**
Right (contact)	0.2025	0.1058
Left (contact)	0.1406	0.0306
Average (contact)	**0.1716**	**0.0682**
Right (all)	0.1982	0.1122
Left (all)	0.1490	0.0316
Average (all)	**0.1736**	**0.0719**

Table 1. RMS Position Error

4. Experiment and Results

The system described above was evaluated by measuring the absolute positions of the feet of pedestrians using data acquired synchronously via motion capture (Vicon Motion Systems). Reflective markers were attached to the walkers' shoes, providing an accurate estimate of 3D foot positions. Five recordings of walking sequences between 5.7 and 12.4 seconds in length were acquired via the apparatus described in Sec. 2. Synchronous motion capture and force data were recorded. Errors were computed based on maximum a posteriori (MAP) foot position estimates obtained from the tracking algorithm.

Figure 7 shows the state estimates and force observations at four stages of the walking sequence: the initial particle set has high variance, but is quickly narrowed down to a few hypotheses which are evolved based on our motion model. Figure 8 shows the resulting motion trajectories and errors for the two foot locations, in planar coordinates. Average RMS error values are reported in Table 1. The experimental parameters are given in Table 2. Position errors during foot-floor contact are found to be slightly less on average than when a foot is not in contact with the floor. Temporal alignment mismatches were found to have a large effect, so we also performed a windowed error calculation in which an acceptable time shift of 10 samples (at 20 Hz) was permitted. This greatly reduced RMS position errors (see Table 1), suggesting that system tracking performance may be most acceptable in situations in which temporal accuracy is not important. Tracking performance in more temporally demanding settings might be greatly improved if a better alignment can be achieved. Video documentation of these results is provided in the supplementary material, and available online here.

In addition to position estimates, this system provides continuous labels identifying the walker's right and left feet. During tracking of the sequences used for evaluation, left and right feet were continuously and coherently identified with 100% accuracy. The capability of this system to maintain and propagate these labels may be useful for appli-

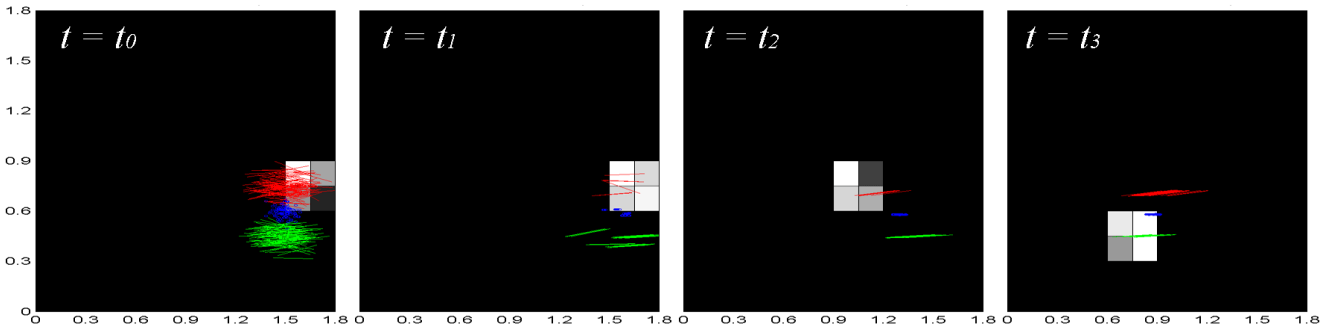

Figure 7. Force observations and pose estimates at 4 stages of the sample walking sequence. Observations are shown as 2D pressure distributions, with quadrant intensity proportional to force sensor readings. Poses are illustrated as green and red lines, corresponding to a top view of left and right feet estimates respectively, with blue circles corresponding to inferred center of mass.

SIR Algorithm	
Number of particles	$N_s = 200$
Initial Noise	$\sigma_S = 4.572$ cm
Roughening Position Noise	$\sigma_x = 0.003048$ cm/rad
Likelihood Model	
Postural prior stance size	$(\mu_\varsigma, \sigma_\varsigma) = (30.48, 30.48)$ cm
Postural prior stance angle	$(\mu_\Theta, \sigma_\Theta) = (0, \pi/2)$ rad
Dynamic Model	
Additive Position Noise	$\sigma_{u,v} = 0.3048$ cm
Additive Angle Noise	$\sigma_\theta = 0.003048$ rad
Longitudinal bias	$\gamma = 0.9$
Velocity mixing rate	$\alpha = 0.55$
Probability of no β change	$\delta = 0.55$

Table 2. Experimental parameters

cations including floor-based touch screen user interfaces, where it may be desirable to assign each foot a different functional operation, or to render a different response to left and right foot [20].

5. Conclusion

This paper presented Bayesian filtering techniques to track the lower body pose of a pedestrian from foot-floor interaction forces acquired via a coarse array of in-floor force sensors. The system achieves continuous and labeled tracking of the lower limbs of a walker via a coarse sensor array, by combining prior knowledge about the mechanical structure of the interface and a simple, but consistent, model of the dynamics of the feet during walking. In the experiment described above, our system never confused the left and right feet of the walker, and was able to track the locations of each with an average resolution on the order of 15 cm, and with improved resolution when feet are in contact with the surface. Potential applications of these techniques include tracking in smart environments in which motion capture is impractical (due to occlusion or other factors), and to interaction with distributed, floor-based touch surface interfaces for the feet [20, 19].

Despite the promising nature of the results presented above, further experiments are needed to evaluate the quality of the tracking. In addition, it is clear that the system itself can be improved in several respects. A higher-density sensor network would improve position estimates during contact, albeit at greater cost. A model for the non-contact portion of walking movement that is more sophisticated than the random drift model used in our system could significantly improve estimates in foot tracking when a foot is not in contact with the floor. The incorporation of additional prior knowledge about the kinematic constraints on lower limb positions during walking would also be expected to contribute improvements. In ongoing work, we are exploring the possibilities for optimally fusing information from in-floor force sensors with motion capture or other video sensors, in order to resolve data loss due to camera occlusion, or to provide contact forces and timing information that cannot be accurately estimated from video.

References

[1] M. Addlesee, A. H. Jones, F. Livesey, and F. S. Samaria. The orl active floor. *IEEE Personal Communications*, 4:35–41, 1997.

[2] A. Bicchi, J. Salisbury, and D. Brock. Contact sensing from force measurements. *The International Journal of Robotics Research*, 12(3):249, 1993.

[3] M. Brubaker, L. Sigal, and D. Fleet. Estimating Contact Dynamics. In *IEEE International Conference on Computer Vision (ICCV)*, 2009.

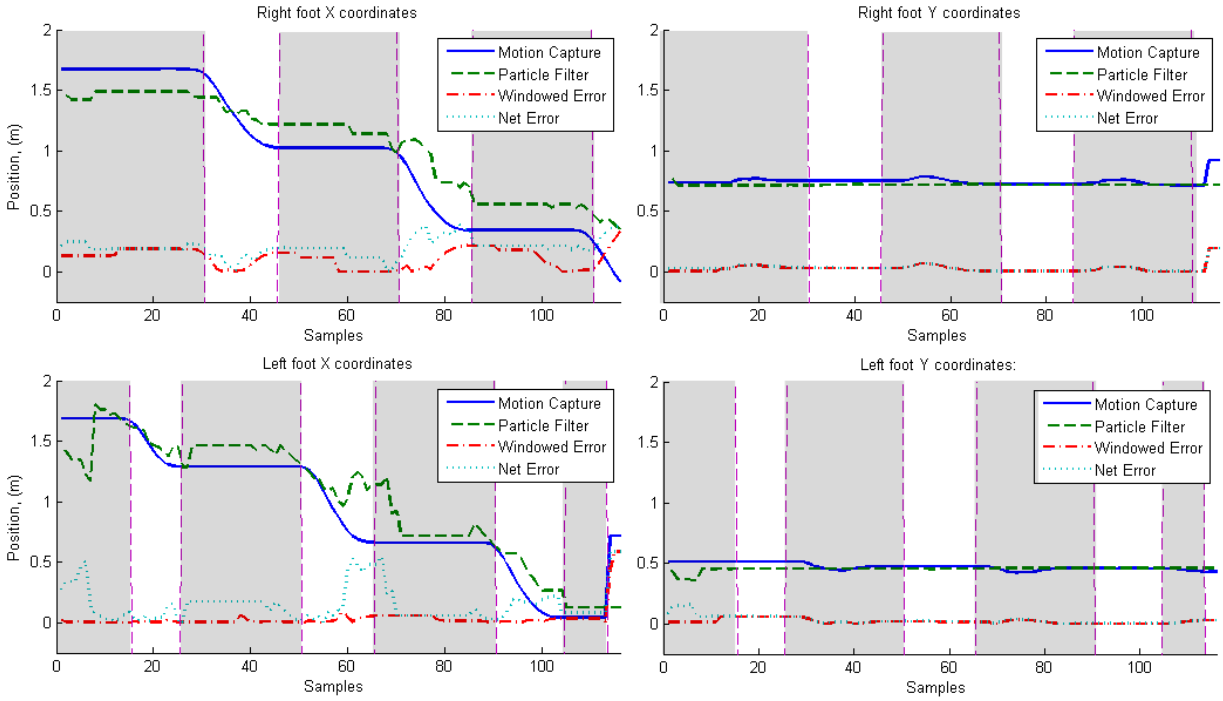

Figure 8. Sample results for a pedestrian crossing the apparatus, comparing measurements from motion capture data and estimates obtained from the particle filter. Results are shown for both feet ϕ_l, ϕ_r. The shaded domains illustrate contact states, with grey and white corresponding to $\beta_i = 1$, $\beta_i = 0$ respectively.

[4] D. Fox, J. Hightower, H. Kauz, L. Liao, and D. Patterson. Bayesian techniques for location estimation. In *Proceedings of The 2003 Workshop on Location-Aware Computing*, page 16, 2003.

[5] V. Fox, J. Hightower, L. Liao, D. Schulz, and G. Borriello. Bayesian filtering for location estimation. *Pervasive Computing, IEEE*, 2(3):24–33, July-Sept. 2003.

[6] L. Middleton, A. Buss, A. Bazin, and M. Nixon. A floor sensor system for gait recognition. In *Proceedings of the Fourth IEEE Workshop on Automatic Identification Advanced Technologies*, pages 171–176, 2005.

[7] T. Moeslund and E. Granum. A survey of computer vision-based human motion capture. *Computer Vision and Image Understanding*, 81(3):231–268, 2001.

[8] T. Moeslund, A. Hilton, and V. Krueger. A survey of advances in vision-based human motion capture and analysis. *Computer vision and image understanding*, 104(2-3):90–126, 2006.

[9] H. Morishita, R. Fukui, and T. Sato. High resolution pressure sensor distributed floor for future human-robot symbiosis environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2002), Lausanne, Switzerland*, pages 1246–1251, 2002.

[10] T. Murakita, T. Ikeda, and H. Ishiguro. Human tracking using floor sensors based on the Markov chain Monte Carlo method. In *Pattern Recognition, 2004. ICPR 2004. Proceedings of the 17th International Conference on*, volume 4, 2004.

[11] R. J. Orr and G. D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 275–276, New York, NY, USA, 2000. ACM.

[12] J. Paradiso, C. Abler, K.-y. Hsiao, and M. Reynolds. The magic carpet: physical sensing for immersive environments. In *CHI '97: CHI '97 extended abstracts on Human factors in computing systems*, pages 277–278, New York, NY, USA, 1997. ACM.

[13] R. F. Pinkston. A touch sensitive dance floor/midi controller. *The Journal of the Acoustical Society of America*, 96(5):3302–3302, 1994.

[14] I. Rius, X. Varona, F. X. Roca, and J. Gonzàlez. Posture constraints for bayesian human motion tracking. In *Articulated Motion and Deformable Objects*, pages 414–423. Springer Verlag, 2006.

[15] A. Schmidt, M. Strohbach, K. v. Laerhoven, A. Friday, and H.-W. Gellersen. Context acquisition based on load sensing. In *UbiComp '02: Proc of the 4th international conference on Ubiquitous Computing*, pages 333–350, London, UK, 2002. Springer-Verlag.

[16] J. Suutala, K. Fujinami, and J. Roening. Gaussian Process Person Identifier Based on Simple Floor Sensors. In *Proceedings of the 3rd European Conference on Smart Sensing and Context*, page 68. Springer-Verlag, 2008.

[17] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.

[18] Y. Visell and J. Cooperstock. Design of a Vibrotactile Display via a Rigid Surface. In *Proc. of IEEE Haptics Symposium*, 2010.

[19] Y. Visell and J. Cooperstock. Visell, Y. and Law, A. and Smith, S. and Ip, J. and Cooperstock, J. In *Proc. of IEEE Haptics Symposium*, 2010.

[20] Y. Visell, A. Law, S. Smith, R. Rajalingham, and J. Cooperstock. Contact sensing and interaction techniques for a distributed multimodal floor display. In *Proc. of IEEE 3DUI*, 2010.

[21] K. Yin and D. K. Pai. Footsee: an interactive animation system. In *SCA '03: Proc of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 329–338, 2003.