

A Context-Based State Estimation Technique for Hybrid Systems

Sarjoun Skaff* Alfred A. Rizzi* Howie Choset* Pei-Chun Lin†
sarjoun@ri.cmu.edu arizzi@ri.cmu.edu choset@ri.cmu.edu pclin@engin.umich.edu

* Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA

† Department of Mechanical Engineering, The University of Michigan, Ann Arbor, MI, USA

Abstract—This paper proposes an approach to robust state estimation for mobile robots with intermittent dynamics. The approach consists of identifying the robot's mode of operation by classifying the output of onboard sensors into mode-specific contexts. The underlying technique seeks to efficiently use available sensor information to enable accurate, high-bandwidth mode identification. Context classification is combined with multiple-model filtering in order to significantly improve the accuracy of state estimates for hybrid systems. This approach is validated in simulation and shown experimentally to produce accurate estimates on a jogging robot using low-cost sensors.

Index Terms—State Estimation, Classification, Multiple-Model Filtering, Hybrid Systems

I. INTRODUCTION

Robust state estimation is a key enabling technology for reactive robotic systems. Mobile robots traversing rough terrain often exhibit complex intermittent dynamics that are difficult to model accurately and thus limit the performance of state estimation filters. Hybrid systems may frequently switch between modes of operation and experience complex transients during mode transitions. Transitions may occur at rates similar to the bandwidth of onboard sensors and cause traditional filters to diverge. These considerations motivate the construction of estimation systems which can efficiently make use of all available sensor data to reliably adapt to the current operating mode and generate accurate state estimates.

This paper presents a framework to take advantage of all measurements to capture as much information as possible about the dynamics causing the robot's observed behavior. Simulations of a bouncing ball show that the framework significantly increases the accuracy of estimating the balls' height. Experiments are also conducted on RHex, a highly mobile six-legged robot which locomotes by rotating its half-circle shaped compliant legs. Careful synchronization of leg rotations produce different tripod gaits, enabling the robot to walk, jog and run [1]. Of particular interest is the estimation of RHex's state while jogging, as the gait produces complex dynamics with alternating flight and stance phases. The challenge is to rely on fairly inexpensive onboard sensors such as accelerometers to accurately estimate the height of the robot in real time.

The approach involves building an information processing system that identifies a robot's current dynamical context.

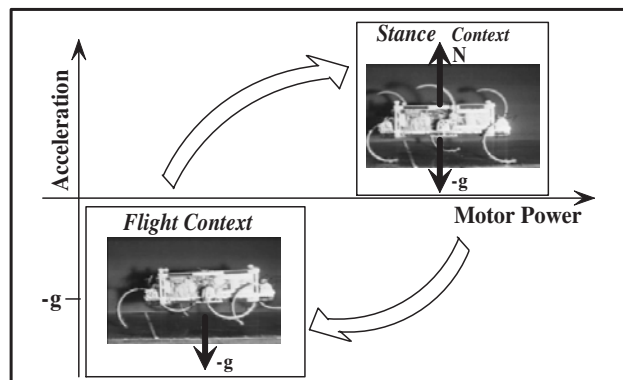


Fig. 1. Flight and stance contexts correspond to data generated by acceleration and motor power sensors while the robot operates in these modes. The flight context corresponds to measurements of flight conditions, such as calibrated accelerations close to $-g$ and low motor power consumption. Conversely, the stance context corresponds to positive accelerations (resultant of ground reaction and gravity) and high power consumption. RHex pictures courtesy of Haldun Komsuoğlu.

Contexts correspond to a classification of data generated by onboard sensors when the robot operates in specific modes. Fig. 1 shows an example of two contexts that correspond to a legged robot that alternates flight and stance modes. The mode of operation can be identified by positively classifying current sensor data in one of the available contexts. The classification seeks to incorporate as much available information as possible in order to identify the modes at a bandwidth comparable to that of onboard sensors. Applied to hybrid systems, the mode identification mechanism can be combined with traditional multiple-model filters to help select accurate models and significantly improve the accuracy of state estimates.

With technological and scientific advances, the locomotion capability and complexity of mobile robots is increasing. Fortunately, robot behavior can often be approximated with simple models that abstract the complex robot-environment interactions [2]. More formally, consider the space of all robot states, \mathcal{Q} , and the space \mathcal{Q}_m of states that the simple model can express. Generally $\dim(\mathcal{Q}_m) < \dim(\mathcal{Q})$, and \mathcal{Q}_m is the image of the map $\pi_Q : \mathcal{Q} \rightarrow \mathcal{Q}_m$.

The sensor space \mathcal{S} is the set of measurements produced by mapping the state $x \in \mathcal{Q}$ to observations $y \in \mathcal{S}$ through

$g : \mathbb{Q} \rightarrow \mathbb{S}$. It is often the case with mobile robots that sensors measure components of the robot state that are not components of the model state. Hence, sensor measurements y are mapped to measurements $y_m \in \mathbb{S}_m$ through $\pi_S : \mathbb{S} \rightarrow \mathbb{S}_m$, where \mathbb{S}_m is the set of measurements that are compatible with the model. In a sense, y_m is a measurement of the model state described by the mapping $h : \mathbb{Q}_m \rightarrow \mathbb{S}_m$. The following commutativity diagram summarizes the relationships among \mathbb{Q} , \mathbb{Q}_m , \mathbb{S} and \mathbb{S}_m :

$$\begin{array}{ccc} \mathbb{Q} & \xrightarrow{g} & \mathbb{S} \\ \downarrow \pi_Q & & \downarrow \pi_S \\ \mathbb{Q}_m & \xrightarrow{h} & \mathbb{S}_m \end{array}$$

II. RELATION TO PREVIOUS WORK

The approach for context-based state estimation uses classification techniques to recognize robot dynamics and multiple-model filtering to estimate the state. These techniques are extended beyond their traditional application domain of machine vision and fault detection to address the intermittent and continuously changing dynamics of mobile robots.

A. Multiple-model filtering

State estimation for hybrid systems is a problem long addressed by the scientific community, with particular emphasis on the fields of aircraft fault detection and radar target tracking (e.g. [3], [4], [5], [6], [7]). Some research in the robotics community is adopting multiple-model approaches for fault detection in mobile robots ([8], [9]). However, little attention has been paid to the specifics of estimating the state of robots with hybrid dynamics.

Conventional multiple model estimators based on the Kalman filter framework such as the Interacting Multiple Models (IMM) update multiple filters in parallel and consolidate their output into a state estimate [10], [11]. These techniques have two disadvantages. The first is the necessity of running multiple filters in order to capture mode transitions. This can be computationally expensive, as the number of filters grows at least quadratically with the number of modes. The second disadvantage is the sub-optimal accuracy of the estimate in the special case where the mode is identified unambiguously. This case calls for only updating the filter corresponding to the identified mode, and not averaging its output with knowingly inaccurate estimates from other filters, as this would lower the quality of the output. However, IMM-like techniques require that the complete bank of filters be updated in order to evaluate the accuracy of the model set. As a consequence, they cannot deactivate unnecessary filters and still decide when to reactivate them.

The proposed approach attempts to remedy these problems. It provides a technique that turns off irrelevant filters and correctly re-initializes them when they become accurate. This saves computational cost and most importantly avoids

sub-optimal estimates when the mode is identified unambiguously.

B. Classification

Context identification through classification is used in the computer vision community to enhance object recognition, and in the robotics community to adapt sensor parameters to the environment, assess the accessibility of a terrain and even build motion models [12], [13], [14], [15].

The essentially geometric classification effective in machine vision problems assumes that changes in signals are slow [12]. Classification of dynamic data is addressed by the robotics community, but under the assumption that the world can be described with a set of discrete, static states [13]. Techniques that do account for continuously changing states assume that the dynamics are time-invariant [15]. In summary, available techniques assume that either no dynamics are involved, or that the dynamics do not change with time.

When applied to field robots, classification cannot rely on these assumptions because robot dynamics continually alter sensor output. Therefore, the advocated approach for context identification seeks to increase the bandwidth of classification by incorporating as much sensor information as possible. In addition, it accounts for time varying dynamics by combining classification and multiple-model filtering to capture dynamics that could no longer be classified with available contexts.

III. CONTEXT-BASED STATE ESTIMATION

The concept of context-based state estimation is to identify robot operating modes by classifying onboard sensor data into contexts that correspond to specific modes of operation. Context classification essentially uses sensor information to build statistical models that recognize system dynamics. Assuming the availability of models that describe the different dynamics of a hybrid system, this information helps estimators choose appropriate models for state estimation.

A naive approach to state estimation would be to attempt to compute the state directly from $\pi_Q^{-1} \circ h^{-1} \circ \pi_S : \mathbb{S} \rightarrow \mathbb{Q}$. However, π_Q and h are generally not invertible, so system state is estimated with filters that process model prediction and sensor information mapped by π_S . Since π_S maps sensor output to measurements compatible with the model, it may discard some information. The simpler the model and the greater the complexity of the dynamics, the less information is available to the filter and the longer it will take to converge. In contrast, context classification does not necessarily depend on complex models and instead seeks to efficiently incorporate all available information. Thus, modes can be identified accurately and at bandwidth comparable to that of onboard sensors. These advantages compel the use of classification rather than multiple-model filtering for mode identification. Context classification provides multiple-model filters with accurate, high-bandwidth

mode identification which helps improve their convergence rate and the accuracy of their state estimates.

A. Context Predicates and Contexts

Sensor output that is classified into contexts is called a *context predicate*, as it helps identify the current context. Context predicates can also be constructed from processed sensor data (e.g. IMU output), state estimates from a filter, or filter components (e.g. residual).

Contexts are defined as the set of all predicate values that correspond to specific modes of operation. Different contexts are assumed to correspond to modes of operation with distinct dynamics, so dynamics are uniquely mapped to their corresponding context under the presumption that the set of predicates is sufficiently rich to ensure that contexts do not overlap.

Contexts are constructed by operating the robot in different modes while collecting predicate values. The collection of predicates corresponding to each mode is a bounded set in \mathbb{S} that constitutes a context. Future predicate values are expected to fall within the context's bounds when the corresponding mode is in operation. Thus, a robot's mode can be identified by classifying current predicates within available contexts.

To illustrate the concept, consider RHex executing a jogging gait, with onboard sensors measuring acceleration and power consumption. During the flight phase, the measured acceleration should be close to gravity ($-g$) with little motor power consumed. During stance, the legs should produce a positive acceleration and consume additional power. This reasoning can be formalized by classifying the acceleration and the power consumption predicates into flight and stance contexts, as shown in Fig. 1. The phase of RHex's gait could be identified by comparing current values of the predicates to the two contexts. Whereas additional contexts that capture modes such as uneven leg touchdown and foot slippage may be able to describe the dynamics more accurately, for clarity this paper only considers two modes. This is sufficient to evaluate the effectiveness of a collection of simple models at capturing complex behavior.

B. Multiple-Model Filtering

As introduced in Section II-A, the leading multiple-model estimation technique is the IMM. Each IMM iteration starts with the assumption that any mode could have been in effect at time $t - 1$, and any mode could be in effect at time t . For a system with N modes, a bank of N^2 filters is updated and the output of all filters is consolidated. Algorithm 1 details the steps involved in an IMM cycle. Here, $(i, j) \in \{\text{set of modes}\}$; $filter_i$ is based on the model of mode i ; i_t represents the hypothesis that mode i is in effect at time t ; $x_{i,j}$ is the state predicted by $filter_i$ and corresponding to the sequence (i_t, j_{t-1}) ; r is the residual; and S and P are the innovation and process covariances,

Algorithm 1 Steps of an IMM update [10]

$$[x_{i,j}, P_{i,j}] = filter_i(x_j) \quad (1)$$

$$p_{i,j} = \frac{1}{\sqrt{2\pi S_{i,j}}} \exp\left(-\frac{1}{2}r_{i,j}S_{i,j}^{-1}r_{i,j}^T\right) \text{ (Likelihood)}$$

$$T_{i,j} = \text{Prob}(i_t|j_{t-1}) \text{ (Transition Probability)} \quad (2)$$

$$Prob_{i,j} = \frac{p_{i,j}T_{i,j}Prob_i}{\sum_j p_{i,j}T_{i,j}Prob_i} \text{ (Probability}(i_t, j_{t-1})) \quad (3)$$

$$x_i = \sum_j x_{i,j}Prob_{i,j} \text{ (Mode State)} \quad (4)$$

$$P_i = \sum_j Prob_{i,j} \left(P_{i,j} + (x_{i,j} - x_i)(x_{i,j} - x_i)^T \right)$$

$$Prob_i = \frac{\sum_j p_{i,j}T_{i,j}Prob_i}{\sum_i \sum_j p_{i,j}T_{i,j}Prob_i} \text{ (Probability}(i_t)) \quad (5)$$

$$x = \sum_i x_i Prob_i \text{ (Consolidated State)} \quad (6)$$

$$P = \sum_i Prob_i \left(P_i + (x_i - x)(x_i - x)^T \right)$$

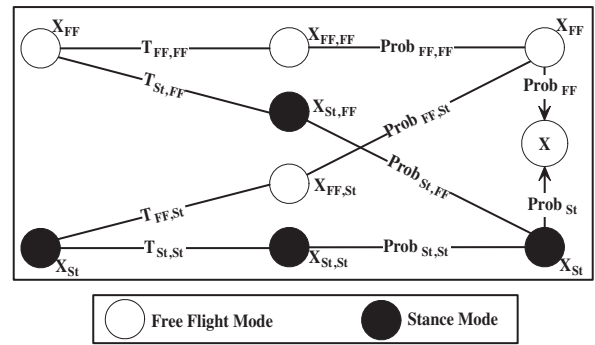


Fig. 2. A two-mode IMM spawns $N^2 = 4$ filters, one for each mode sequence. Their output is consolidated to form the state estimate.

respectively. Fig. 2 illustrates the cycle for a two-mode system (FF = Free Flight; St = Stance).

When the context predicates fall within the bounds of a context, the corresponding mode is identified unambiguously and filters should rely on the corresponding model to estimate the state, at the exclusion of other inaccurate models. However, IMM does not discard knowingly inaccurate estimates, which leads to suboptimal consolidated states.

Two modifications of the IMM alleviate this problem. The first is to change the transition probabilities, $T_{i,j}$, in (2) as a function of the mode. If mode i is recognized as being in operation, and j represents all other modes, then set $T_{i,i} = T_{i,i} = 1$ and $T_{j,i} = T_{j,j} = 0$. This means that transitioning into the identified mode and staying in it has a probability of one, and transitioning into a wrong mode and staying in it has a probability of zero. As expected, this produces $Prob_i = 1$ and $Prob_j = 0$.

The second modification is to set $x = x_i = x_{i,i}$. This formalizes the observation that once the system is in the mode, it is expected to remain in it until a change of context is detected. In other words, the only valid hypothesis is (i_t, i_{t-1}) , and since different modes correspond to distinct dynamics, the output of $filter_i(x_i)$ in (1) is not combined to other estimates and instead constitutes the sole output of the IMM. By ignoring the contribution of inaccurate mode states, the accuracy of the consolidated state is not decreased unnecessarily. In addition, ignored filters can now be deactivated with no impact on state estimates. For example, when the two-mode system of Fig. 2 is in flight, only $x_{FF,FF}$ is updated while the other filters are deactivated.

When the context can no longer be identified unambiguously, all filters can be restarted and the IMM resumes nominal operation. The output of the filters are consolidated (Equations 4 and 6), the mode probabilities are computed (Equations 3 and 5), and the mode transition point estimated (when the probability of one mode exceeds that of all others). Thus, context-based filtering combines the advantage of two techniques. Classification provides accurate mode identification when it is applicable, and multiple-model filtering estimates mode transition points when no other identification means are available.

IV. SIMULATION RESULTS

As a first step towards robot deployment, context-based state estimation is implemented on a simulated hybrid system. The setup consists of an elastic bouncing ball that alternates flight and stance dynamics:

$$\ddot{x} = \begin{cases} -g & \text{Ballistic Projectile} \\ K(x - l_0) - g & \text{Loss-less Mass-Spring System} \end{cases}$$


where the state x is height, g is the acceleration due to gravity and K and l_0 are the ball's effective spring constant and rest length, respectively. The task is to estimate the ball's height and velocity using the flight and stance models and a noisy height sensor.

A. Conventional IMM


The initial implementation deploys the conventional IMM of Fig. 2. For lack of better information, $T_{i,j} = \frac{1}{2}, \forall i, j \in \{FF, St\}$. Four filters are updated for this two-mode system, one per mode sequence (i_t, j_{t-1}) . Equation 1 outputs four states $x_{i,j}$, which are consolidated through (4) and (6) into x , the ball state estimate (Fig. 3). Over 10s of simulation and one complete bounce cycle, the RMS error for position estimates is 6.34 cm and for velocity estimates is 15.5058 m/s. The poor accuracy results from incorporating inaccurate stance estimates during flight and flight estimates during stance.

B. Context-Based IMM


To address this problem, a flight and a stance contexts are introduced. The predicate is the height of the ball as estimated by the IMM. This choice is motivated by the



(a) Velocity estimates



(b) Error in position and velocity estimates



(c) Mode Probabilities

Fig. 3. State estimates obtained with a conventional IMM lead to RMS errors of 6.34 cm and 15.5058 m/s for position and velocity estimates, respectively. Modifications to the IMM discard knowingly inaccurate estimates and reduce RMS errors to 0.13 cm and 0.0119 m/s.

observation that when the ball is at high altitude, it is known to be in flight, and when it is compressed, it is known to be in stance. Thus, classifying the ball's height as high (flight) or low (stance) is a rational approach at recognizing the dynamics.

The classification fails in the immediate vicinity of touch-down, $x = l_0 \pm \delta$, with $\delta \ll l_0$, where it is unclear when the ball transitions from one mode to the other. Therefore, the contexts are defined as follows:

- Flight Context: $x > l_0 + \delta$
- Stance Context: $x < l_0 - \delta$
- Unidentified Context: $l_0 - \delta < x < l_0 + \delta$

Pseudo code for the modified IMM is provided in Algorithm 2. When the ball is recognized to be in flight ($x > l_0 + \delta$) or in stance ($x < l_0 - \delta$), only the corresponding $filter_{FF}(x_{FF})$ or $filter_{St}(x_{St})$ is updated. When

Algorithm 2 IMM Modifications

<pre> for $i \in \{FF, St\}$, $j = \bar{i}$ if (identified mode = i) $T_{i,j} = T_{i,i} = 1$ $T_{j,i} = T_{j,j} = 0$ $x = x_i = x_{ii}$ $flag_i = T$, $flag_j = F$ if(no mode identified){ for $i \in \{FF, St\}$, $j = \bar{i}$ if{$flag_i = F$}{ $x_i = x_j = x_{ii}$ $P_i = P_j = P_{ii}$ } $flag_{FF} = flag_{St} = T$ if($\dot{x} < 0$){$i = St$; $j = FF$} else{$i = FF$; $j = St$} $T_{i,j} = 0.7$; $T_{j,j} = 0.3$ $T_{i,i} = 1$; $T_{j,i} = 0$ </pre>	<pre> for $i \in \{FF, St\}$, $j = \bar{i}$ if($flag_i$){ update $filter_i(x_i)$ if($flag_j$){ update $filter_i(x_j)$ } } if($flag_j$){ update $filter_j(x_j)$ if($flag_i$){ update $filter_j(x_i)$ } } </pre>
---	---

the context is unidentified, all four filters are updated. The transition probabilities in (2) are adjusted so that when the ball is descending ($\dot{x} < 0$), the probability of transitioning into stance is greater than the probability of staying in flight (e.g. $T_{St,FF} = 0.7$; $T_{FF,FF} = 0.3$). The inverse is true when the ball is ascending ($T_{FF,St} = 0.7$; $T_{St,St} = 0.3$). Fig. 3(a) shows that this strategy significantly improves the accuracy of the estimated velocity. RMS errors are reduced to 0.13 cm and 0.0119 m/s for position and velocity, respectively. Fig. 3(b) shows that estimate errors are small until touchdown and rise only modestly thereafter, in sharp contrast with the estimate errors of the conventional approach. This illustrates the benefit of discarding knowingly inaccurate estimates. Fig. 3(c) shows the selective deactivation of unnecessary filters during flight and stance as well as the transition point from one mode to the other.

V. EXPERIMENTATION RESULTS

This technique is extended to RHex, where the task is to estimate the jogging robot’s height using onboard accelerometers. RHex’s locomotion mechanism is characterized by six compliant legs that confer to the robot a significant capacity to overcome obstacles, but also generate large dynamics that are difficult to model [16], [2].

A. State Estimation With Leg Strain Gauges

Fig. 4 shows RHex’s vertical acceleration when jogging. The frequency of the signal corresponds to that of the gait alternating flight and stance phases. The identity of the mode in operation can be inferred from strain gauges that measure leg compression and thus detect whether the robot is in flight or in stance [17]. Therefore, strain gauge measurements serve as a predicate for the flight and stance contexts.

For simplicity, RHex is modeled as a bouncing ball, using the ballistic flight and elastic mass-spring models. The




Fig. 4. Rhex’s acceleration in the vertical direction as recorded by on-board accelerometer. The gravity bias is removed from the measurement, so a zero acceleration corresponds to the robot standing still. Leg strain gauges detect stance phases (diamonds) and flight phases (circles).

mass-spring parameters are set to the actual mass of the robot (8.5 Kg) and the spring constant that yields acceptable state estimates is found to be 6800 N.m, a value close to the physical spring constant estimated at 6600 N.m.

Kalman filters combine predicted and measured accelerations and output consolidated estimates. When the strain gauge measurements indicate that the robot is in stance or in flight, the IMM only updates the corresponding filter. Transitions between flight and stance compress the legs in complex sequences that make context identification difficult. In these situations, the complete set of IMM models are updated. The output is then integrated twice to extract the state.

The resulting height estimates compare favorably with ground truth, as shown in Fig. 5(a). The mean and standard deviation of the RMS difference between the estimated height and the ground truth measurement over 12 experiments is 1.85 cm and 0.37 cm, respectively. Fig. 5(b) shows the system’s estimate of mode probabilities, the transition points from one mode to the other, and the selective activation of the filters.

The quality of the estimates underlines the significance of the mass-spring model’s contribution. Through the virtual spring, the model anchors the robot to the ground and maintains its height within the vicinity of l_0 . Thus, by expressing the fact that the robot cannot leave the ground, the simple model is able to significantly mitigate the problem of filter divergence. For comparison, height estimates obtained by directly integrating the accelerometer are plotted. As expected, the model-less estimates diverge rapidly.

B. State Estimation Without Leg Strain Gauges

Strain gauges enable accurate mode detection, but they are expensive and complex to install on circulating legs. Accelerometers may provide a less expensive alternative. This can be achieved by reasoning about what acceleration values are expected in each mode. For instance, calibrated accelerations close to $-g$ are expected during flight, and positive accelerations are expected during stance. Thus, instead of identifying the mode by classifying strain gauge measurements, the strategy is to classify accelerometer output as follows:

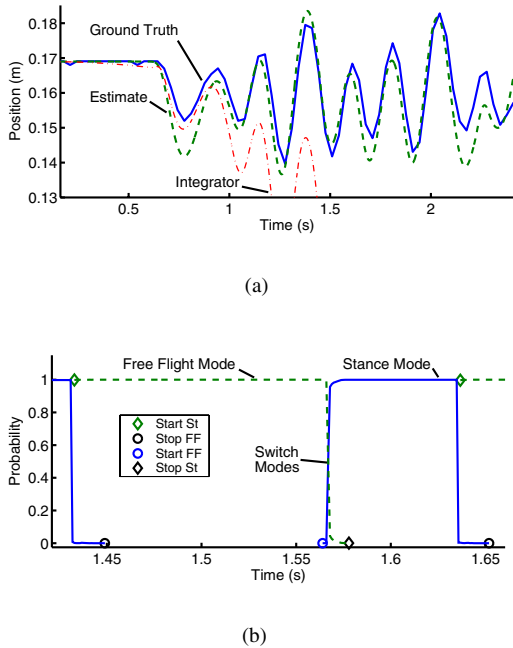


Fig. 5. The modified IMM yields relatively accurate estimates of RHex's height while jogging. Here, the robot's behavior is modeled after an elastic bouncing ball.

- Flight Context: $\ddot{x} < -6 \text{ m/s}^2$
- Stance Context: $\ddot{x} > 0$
- Unidentified Context: $0 > \ddot{x} > -6 \text{ m/s}^2$. Accelerometer output within this context indicates that the robot is transitioning between flight and stance. Context bounds are derived from experiment observation.

State estimates obtained with this contextual classification are virtually indistinguishable from the results of Fig. 5. The mean and standard deviation of the RMS error over 12 experiments are 2.09 cm and 0.39 cm, respectively, which compare favorably to the values obtained with the help of strain gauges. This result indicates that accurate state estimation could be performed using low-cost sensors such as accelerometers.

VI. CONCLUSION

Context-based state estimation enables accurate, high-bandwidth mode identification for hybrid systems. The approach classifies sensor measurements in mode-specific contexts to recognize the dynamics and help multiple-model filters choose appropriate models.

This research extends the applicability of classification to hybrid systems with time varying dynamics and introduces modifications to the IMM algorithm in order to take advantage of independent mode identification. When classification recognizes the dynamics, the mode is identified at a bandwidth similar to that of onboard sensors. Modified IMMs use this information to selectively update appropriate filters and deactivate others. When the changing dynamics

cannot be classified, IMMs resume conventional operation for estimating state and mode transition points, which in turn helps improve context identification. Thus, combining context classification and multiple-model filtering improves the bandwidth and accuracy of state estimation.

Ongoing research is extending context classification to the problem of verifying the accuracy of the entire model set used by a multiple-model system. The technique is also being applied to systems with smoothly varying dynamics, where the challenge is to estimate the accuracy with which the dynamics are recognized, infer an assessment of model accuracy, and adapt filter parameters accordingly.

ACKNOWLEDGEMENT

This work is supported by DARPA/ONR Contract N00014-98-1-0747 and DARPA/SPAWAR N66001-03-C-8045.

REFERENCES

- [1] U. Saranli, M. Buehler, and D. E. Koditschek, "RHex: A Simple and Highly Mobile Robot," *International Journal of Robotics Research*, vol. 20, pp. 616–631, July 2001.
- [2] S. Skaff, G. Kantor, D. Maiwand, and A. A. Rizzi, "Inertial navigation and visual line following for a dynamical hexapod robot," in *2003 IEEE/RSJ International Conference on Intelligent Robots and Systems*, (Las Vegas, NV), October 2003.
- [3] E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan, "Interacting multiple model methods in target tracking: a survey," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, 1998.
- [4] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Mathematics in science and engineering, Academic Press, 1988.
- [5] Y. Zhang and J. Jiang, "Integrated active fault-tolerant control using IMM approach," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, pp. 1221–1235, October 2001.
- [6] P. S. Maybeck and P. D. Hanlon, "Performance enhancement of a multiple model adaptive estimator," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 31, October 1995.
- [7] K. A. Fisher and P. S. Maybeck, "Multiple model adaptive estimation with filter spawning," in *Proceedings of the American Control Conference*, (Chicago, Illinois), pp. 2326–2331, June 2000.
- [8] V. Verma, J. Fernandez, and R. Simmons, "Probabilistic models for monitoring and fault diagnosis," in *The Second IARP and IEEE/RAS Joint Workshop on Technical Challenges for Dependable Robots in Human Environments*. (R. Chatila, ed.), October 2002.
- [9] S. Thrun, J. Langford, and V. Verma, "Risk sensitive particle filters," December 2001.
- [10] C. B. Chang and M. Athans, "State estimation for discrete systems with switching parameters," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 14, no. 3, pp. 418–425, 1978.
- [11] P. S. Maybeck, *Stochastic Models, Estimation and Control*, vol. 1 of *Mathematics in science and engineering*. Academic Press, 1979.
- [12] A. Torralba, K. Murphy, W. Freeman, and M. Rubin, "Context-based vision system for place and object recognition," in *The 9th International Conference on Computer Vision*, (Nice, France), 2003.
- [13] S. Lenser and M. Veloso, "Automatic detection and response to environmental change," in *2003 International Conference on Robotics and Automation*, (Taipei), September 2003.
- [14] C. Wellington and A. Stentz, "Learning predictions of the load-bearing surface for autonomous rough-terrain navigation in vegetations," in *Proceedings of the International Conference on Field and Service Robotics*, pp. 49–54, July 2003.
- [15] D. Nikovski, *State-Aggregation Algorithms for Learning Probabilistic Models for Robot Control*. PhD thesis, Carnegie Mellon University, 2002.
- [16] U. Saranli, *Dynamic Locomotion with a Hexapod Robot*. PhD thesis, University of Michigan, September 2002.
- [17] P.-C. Lin, H. Komsuoğlu, and D. E. Koditschek, "A leg configuration sensory system for dynamical body state estimates in a hexapod robot," in *IEEE Int. Conf. Robotics and Automation*, 2003.