

Low-cost Accelerometers for Robotic Manipulator Perception

Morgan Quigley, Reuben Brewer, Sai P. Soundararaj, Vijay Pradeep, Quoc Le, and Andrew Y. Ng

Abstract— We present a series of experiments which explore the use of consumer-grade accelerometers as joint position sensors for robotic manipulators. We show that 6- and 7-dof joint angle estimation is possible by using one 3-d accelerometer for each pair of joints. We demonstrate two calibration approaches and experimental results using accelerometer-based control in both position-control and torque-control regimes. We present a manipulator design combining accelerometer-based sensing with low-cost actuation, and conclude by demonstrating the utility of consumer-grade accelerometers even on high-precision manipulators.

I. INTRODUCTION

Many robotic applications demand manipulators with high degrees of repeatability, precision, and reliability. Production-line robots provide stereotypical examples of this domain. However, although high-precision manipulators excel at tasks for which they can be pre-programmed, robotic manipulators have not become commonplace in the unstructured domains of typical homes and workplaces.

Prior work has investigated the manipulation of objects without high-precision knowledge of its exact position, orientation, and shape [1]. Uncertainty must be resolved by the robot's sensors, such as cameras, lasers, or contact sensors, among other modalities [2]. Furthermore, in some such environments the task itself is not clearly defined, and must be inferred by the robot through human-robot interaction, scene inference, or other methods [3]. This body of research demonstrates that the *effective* manipulation error in unstructured environments is no longer restricted to the mechanical precision of the manipulator; rather, it incorporates errors from the entire robotic system. Sensor calibration, object localization algorithms, and sensor-manipulator registration, among other sources of error, are cumulative and add up quickly. For many tasks envisioned for household service robots, millimeter-order repeatability errors are unlikely to cause drastic changes in the performance of the system, as errors of that order are often overshadowed by larger errors in the perception system, and thus the robotic system must already be robust to such perturbations.

In this paper, we present a sensing strategy which leverages the recent progress in low-cost 3-d MEMS accelerometers. Our approach mounts at least one 3-d accelerometer for each pair of joints and infers the joint angles using an




Fig. 1. Two manipulators used to demonstrate the utility of accelerometer-based sensing: a Willow Garage PR2 Alpha and a prototype low-cost manipulator.

Extended Kalman Filter (EKF). Because of its low cost, this system of joint-position estimation can be used to augment an existing robotic sensor suite, or to function as the primary sensor for a low-cost robotic manipulator.

II. RELATED WORK

Inertial sensing has been used extensively in recent years for human motion capture. Several companies provide small, lightweight, and networkable inertial units which integrate accelerometers, gyroscopes, and magnetometers, and can be easily attached to human limbs and torsos. These units can be worn for use in film or video-game character modeling [4], virtual reality [5], [6], and activity recognition [7].

Accelerometers are used in virtually every aerial and underwater robot as part of the attitude determination system. However, accelerometers have found other uses in robot navigation: for example, when strapped to legged robots, they can be used to classify the surface, which can help in gait selection and tuning [8].

Prior work using accelerometers in robotic manipulation includes simulation results of configuration estimation [9], employing strapdown inertial units on heavy equipment [10], kinematic calibration [11], and the creation of a fault-detection system [12]. In [13], accelerometers were doubly-integrated using an Extended Kalman Filter (EKF) to estimate the configuration of a SCARA-type manipulator undergoing high-dynamic motions. The use of accelerometers in flexible-link robots was proposed in [14]. A human-robot system was proposed in [15] which incorporated the attachment of accelerometers and other sensors to a human tele-operator.

These prior studies have demonstrated that accelerometers offer a compelling source of information. However, we

M. Quigley, S. P. Soundararaj, Q. Le, and A. Y. Ng are with the Computer Science Department, Stanford University, Stanford, CA, USA {mquigley, saip, quocle, ang}@cs.stanford.edu

R. Brewer is with the Department of Mechanical Engineering, Stanford University, Stanford, CA, USA rdbrewer@stanford.edu

V. Pradeep is with Willow Garage, Inc., Menlo Park, CA, USA vpradeep@willowgarage.com

are not aware of published results of robotic manipulators which emphasize the potential for cost-reduction and higher reliability by using accelerometers for state estimation and control.

III. STATE ESTIMATION

Virtually all robotic manipulators use shaft encoders of some variety (optical, magnetic, resistive, etc.) to determine the kinematic configuration of the manipulator. In this section, we depart from this convention and discuss a sensing scheme based solely on 3-d MEMS accelerometers.

In static conditions, a 3-axis accelerometer returns a 3-d vector pointed at the center of the earth. Since the length of this vector is fixed at 1 g, a 3-axis accelerometer only has two degrees of freedom. Hence, one accelerometer is required for every two rotary joints in a robotic manipulator. However, if the kinematics of the structure allows it, incorporating one accelerometer per rotary joint will increase the robustness of the calibration and the accuracy of the joint-state estimates. This effect will be discussed in the following sections.

A. Joint-angle estimation

In this section, we will describe how we produce a coherent estimate of the manipulator state using accelerometers attached to each link of the kinematic chain. This is a sensor-fusion problem: each accelerometer, by itself, can only determine the direction of a downward-facing vector. However, by using *a priori* knowledge of the kinematic constraints of the manipulator, it is possible to produce a unified state estimate. In our implementation, this is done by an Extended Kalman Filter (EKF). An augmented forward kinematics function is used to predict the accelerometer readings given the current belief state. The EKF algorithm then updates the belief state after observing the true measurement. The following paragraphs will discuss this process in more detail.

First, we note that given a 3-d accelerometer reading and knowledge of the axis of rotation, it is possible to compute the joint angle in all cases except a singularity where the axis of rotation is parallel to vertical. In an all-accelerometer sensing scheme, this set of configurations must be avoided, and the severity of this limitation is dependent on the kinematics of the manipulator and the required task: for example, accelerometer-only sensing will completely fail on SCARA manipulators. However, for the anthropomorphic manipulators considered in this paper, vertical joint configurations are readily avoidable. Furthermore, excursions through vertical joint orientations could be gracefully handled by augmenting the accelerometer measurements with magnetometers, back-EMF sensing, angular rate sensors, or conventional shaft encoders in the EKF framework.

Numerous strategies can be employed to infer the manipulator configuration from accelerometer readings. We chose to use an EKF due to its relatively simple implementation and fast on-line performance. A detailed discussion of the EKF algorithm is beyond the scope of this paper and is presented in many excellent texts [16].

To encourage smooth estimates of the joint position, we define the EKF state space to include the joint angles, velocities, and accelerations:

$$\mathbf{x} = \begin{bmatrix} \theta \\ \dot{\theta} \\ \ddot{\theta} \end{bmatrix} \quad (1)$$

where θ represents the joint angles of the manipulator. The state (plant) update function implements numerical integration and assumes constant acceleration:

$$f(\mathbf{x}) = \begin{bmatrix} \theta + \Delta t \dot{\theta} \\ \dot{\theta} + \Delta t \ddot{\theta} \\ \ddot{\theta} \end{bmatrix} \quad (2)$$

The measurement function needs to predict sensor measurements \mathbf{z} given a state \mathbf{x} . As is often the case in EKF implementations, the measurement function could be made arbitrarily complex to capture more and more properties of the system. However, we found that a measurement function which only predicts the acceleration due to gravity was sufficient to handle the low-frequency regime of our manipulator. Adding additional terms to capture centripetal accelerations and linear accelerations induced by joint-angle accelerations did not change the performance, as will be discussed in Section V.

To predict the measurement \mathbf{z}_i of a particular 3-d accelerometer α_i given the state \mathbf{x} ,

$$\mathbf{z}_i = R_i^{\alpha_i} R_{i-1}^i(\mathbf{x}) \cdots R_0^1(\mathbf{x}) R_w^0 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad (3)$$

where $R_i^{\alpha_i}$ represents the rotation from the accelerometer frame to the frame attached to link i , the rotations $R_{i-1}^i(\mathbf{x})$ are the rotations between the link frames, and R_w^0 is the rotation from the base of the manipulator to the world. The following paragraphs describe these rotations in more detail.

$R_i^{\alpha_i}$ is determined by how accelerometer α_i is physically mounted on link α_i . For our manipulator, it is a permutation matrix followed by a rotation of a few degrees. $R_i^{\alpha_i}$ is a static parameter that can be estimated by the calibration process described in the next section.

$R_{i-1}^i(\mathbf{x})$ is determined by the axial orientation of link i and link $i-1$, as well as by the joint angle θ_i present in the state \mathbf{x} . The link twist can be statically estimated during calibration, but the joint angle must be recursively estimated by the EKF.

R_w^0 is the rotation between the base of the manipulator and the gravitational frame. If the manipulator is stationary, this rotation is constant and can be estimated by static calibration.¹

¹There are numerous situations where the R_w^0 rotation is not constant and needs to be estimated, but they lie in more extreme domains of robotics: manipulation on vehicles traveling on mountainous terrain, on spacecraft, or aboard ships in rough seas, for example. These situations are far beyond the scope of this paper but could be readily addressed by estimating R_w^0 through an inertial unit or some other means.




Fig. 2. Left: an unpowered arm used to evaluate the calibration potential of the accelerometer-based sensing approach. Right: touching the end effector to points on the calibration board.

Using the state-update and measurement functions, the EKF algorithm produces a recursive estimate of the mean and covariance of the state as more timesteps and observations are experienced.

B. Calibration

For a variety of reasons, the accelerometer axes of sensitivity are unlikely to be exactly aligned with the axes of the reference frames attached to manipulator links. Internal misalignments of MEMS accelerometer axes, misalignments incurred when mounting accelerometer-equipped circuit boards to the manipulator links, and manufacturing tolerances of the links themselves combine to result in rotations of several degrees between the axes of the accelerometers and the axes of the kinematic frames. Fortunately, these mechanical imperfections are static, and can be calibrated out.

To demonstrate calibration for low-cost manipulators using accelerometers as the primary sensors, we constructed the unactuated arm shown in Figure 2. This 4-dof mockup has a roughly spherical shoulder and an elbow joint, with link lengths similar to the manipulator shown in Figure 1.

Our calibration scheme for the low-cost prototype is loosely derived from the checkerboard-based calibration method widely used in computer vision [17]. A planar calibration board was placed in the workspace of the manipulator. The end effector of the manipulator was touched to its corners, each time recording the corresponding accelerometer readings. The board was then moved and rotated to a different position and orientation, and the process was repeated to collect a dataset of 20 different measurements covering the manipulator workspace. This provides the scale and skew constraints. This data was augmented by collecting many accelerometer readings of manipulator configurations where the end effector was in contact with a large planar surface such as a tabletop. This serves to provide a planarity constraint.

Formally, the optimization problem can be written as:

$$\arg \min_{L, R} g_1(\alpha, L, R) + \lambda_2 g_2(\alpha, L, R) + \lambda_3 g_3(R) \quad (4)$$

where α are the accelerometer readings in the test set, L represents the estimated link parameters, and R represents the rotation matrices representing the misalignment of each accelerometer frame to its respective link frame.

The first term enforces the known scale of the calibration board.

$$g_1(\alpha, L, R) = \sum_{i,j} \|d_{ij} - \hat{d}_{ij}\| \quad (5)$$

where \hat{d}_{ij} is the distance between the estimated positions of the manipulator:

$$\hat{d}_{ij} = |FK(\alpha_i, L, R) - FK(\alpha_j, L, R)| \quad (6)$$

Here, FK means to run the forward-kinematics joint-angle estimation algorithm described in the previous section to produce estimates of the end-effector positions using the link parameters L . The subscripts i and j identify times in the training set which were gathered from the same position and orientation of the calibration square, which therefore correspond to end-effector positions whose ground-truth distance d_{ij} is known.

The second term of the calibration optimization function, g_2 , corresponds to the planarity constraint of the end-effector positions gathered from the tabletop. Let P be the plane fitted by taking the vectors corresponding to the top two singular values of $Y^T Y$, where Y is the $n \times 3$ matrix whose rows y_i consist of the end-effector positions calculated by the estimated calibration. The sum of the distances between the end-effector positions and their projections onto the fitted plane provides an additional source of information about the severity of the miscalibration. Then,

$$g_2(\alpha, L, R) = \sum_{y_i} \|y_i - proj_P(y_i)\| \quad (7)$$

The final term of the optimization function g_3 encourages the misalignment rotation matrices R_i to remain orthonormal during the optimization:

$$g_3(R) = \sum_i \|R_i^T R_i - I_3\| \quad (8)$$

To calibrate accelerometer-based sensing on a high-precision robot, we used the Willow Garage PR2 Alpha platform as a test bed. This robot has 7-dof manipulators equipped with accurate optical shaft encoders in each joint. The manipulator is well calibrated and the link parameters are known *a priori*. The calibration task is simplified, leaving only the rotations R_i of the accelerometer mounting on the kinematic frame to be determined. We treat the joint angles from the shaft encoders as ground truth and compare these with the angle estimates from the accelerometer readings to calibrate out the rotations. The optimization problem presented below is similar to Equation 4

$$\arg \min_R \sum_i \|\theta_i - \hat{\theta}_i(\alpha, R_i)\| + \lambda \sum_i \|R_i^T R_i - I_3\| \quad (9)$$




Fig. 3. Hold-out test set error during the optimization convergence on prototype manipulator. The horizontal axis shows the iteration number, and the vertical axis shows the mean of the miscalibrations. The numerical optimization drives the average error from 11mm to 2mm.

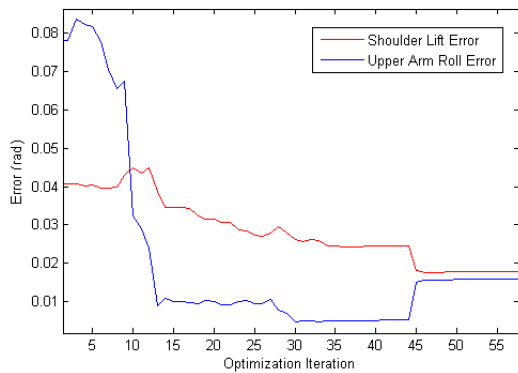


Fig. 4. Hold-out test set error during the optimization convergence on Willow Garage PR2 Alpha. The horizontal axis shows the iteration number, and the vertical axis shows the mean error in joint angle estimates of the shoulder lift and the upper arm roll. The optimization drives the average error from 0.1 radians to 0.02 radians.

where θ_i is the joint angle position of the manipulator as given by the shaft encoders and $\hat{\theta}_i$ is the estimate based on accelerometer readings. $\hat{\theta}_i$ is computed by solving for joint angles by considering inverse kinematics on pairs of links.

Our implementation requires approximately one hour of CPU time to reach convergence using the simplex-based optimization technique implemented by the MATLAB `fminsearch` function. As is always the case with non-convex optimization, a good starting point is necessary to achieve a reasonable solution. We found that starting with the parameters from the CAD model and assuming perfect sensor alignment resulted in reasonable solutions for our test data.

To evaluate the performance of the calibration methods quantitatively, we allowed the optimization algorithm to use training data containing several manipulator positions and checkerboard orientations, and maintained a hold-out test set of several other orientations for evaluation purposes. The results demonstrate a significant calibration accuracy improvement compared to initial guess. Figure 3 shows the convergence of the algorithm from an initial end-effector




Fig. 5. The shoulder motors are evenly spaced around the center of motion to reduce their inertial contribution. The remote center of motion allows the large second and third motors to be direct-drive. The wrist differential is friction-driven by rubber wheels under high compression. This mechanism provides zero-backlash pitch and roll motion.

mean error of 11mm down to 2mm on our prototype manipulator. Figure 4 shows convergence of the algorithm from a mean joint-estimate error of 0.1 radians (5.7 deg) in the shoulder lift to 0.02 radians (1.15 deg) on an instrumented Willow Garage PR2 Alpha.

We note that 2mm mean end-effector localization error is an order of magnitude worse than what is reported by manufacturers of high-quality robotic manipulators sensed by shaft encoders. However, it is more accurate than the best camera-manipulator calibration we have been able to achieve in several years of constructing and calibrating vision-guided systems with high-performance manipulators. We speculate that this level of calibration error would not be the limiting factor of using low-cost localization approaches in a complete robotic system.

IV. A LOW-COST MANIPULATOR

To explore the feasibility of low-cost manipulation using a purely accelerometer-based control scheme, we used a 6-dof manipulator that was constructed in our laboratory under a budgetary constraint of \$1000 USD (Figures 1, 5, 6). As this manipulator incorporates several unconventional design features, we will summarize its design in this section.

The shoulder operates in a spherical RRR configuration with a remote center of motion. Although the shoulder motors are powerful, low-cost, and low-backlash (in the gearhead), they exhibit some cogging due to their ferrous core, which affects the controllability of these joints. Unfortunately, powerful low-cost motors suffer almost uniformly from cogging, and this is a tradeoff that must be balanced in designing and controlling a low-cost arm.

A friction differential drive (Figure 5) provides the wrist with pitch and roll degrees of freedom. The differential is formed by belt-driven rubber wheels pressed firmly against a thin aluminum veneer. Such wheels are low-cost, durable, and effective at transmitting high torques. The friction drive provides zero-backlash and an inherent safety limit on the wrist: overloading results in slippage rather than damage to the drivetrain.

The gripper has two independently-driven fingers, each consisting of a 4-bar linkage for parallel gripping (Figure 6).




Fig. 6. The elbow is driven via belt from a large motor in the shoulder. The gripper is fabricated from lasercut polypropylene, using flexures to create a durable 4-bar mechanism. The thin belts are used only to turn potentiometers for position feedback.




Fig. 7. Accelerometers were attached to the upper arm, forearm, and gripper of a PR2 Alpha robot.

Linkages are constructed from lasercut polypropylene flexures to avoid the mechanical complexity of using discrete parts. Flexures have the further advantages of zero backlash, extreme durability/fatigue resistance, and low-cost.

Linux-based software was written using the open-source Robot Operating System (ROS) platform [18]. Separate programs were written for firmware communication (via the Linux usb-serial driver), state estimation, joint-space control, teleoperation via joysticks, trajectory recording, and trajectory playback/looping. The ROS framework handles the peer-to-peer connections and provides logging, playback, and visualization tools.

V. EXPERIMENTS

In this section, we present a series of experiments that quantify the performance of accelerometer-based state estimation on two robots: the low-cost manipulator discussed in the previous section, and a Willow Garage PR2 Alpha, a high-quality 7-dof manipulator. Accelerometers were designed into the motor control boards of the low-cost manipulator, which were rigidly attached to four links. The PR2 Alpha was outfitted with strapdown accelerometers, as shown in Figure 7.

A. State Estimation

To quantify the performance of the calibration and joint-tracking systems, we affixed accelerometers to a Willow Garage PR2 Alpha robot (Figure 7). This 7-dof manipulator is equipped with high-quality shaft encoders, which served as




Fig. 8. Tracking the forearm roll of the robot shown in Figure 7, showing the encoder ground-truth (red) against the joint angle estimate from the accelerometers (blue).

ground truth for this experiment. Its kinematic configuration includes a vertical first joint (the “shoulder pan”), which we did not attempt to estimate in these experiments, since this joint axis is always parallel to gravity.

Figure 8 demonstrates the tracking performance of one joint (the forearm roll) as the manipulator was smoothly moved through its workspace. The following table (in degrees) shows the mean error throughout the trajectory, measured as the difference between the shaft encoder readings and the joint state estimates from the accelerometers.

	Shoulder Lift	Upper Arm Roll	Elbow
Error (deg)	0.965	0.926	1.590

This experiment was done under near-ideal conditions: the PR2 Alpha arm uses spring balances for gravity compensation and small coreless motors [19]. The mechanism is thus extremely smooth and well-behaved, avoiding any transients or other anomalies as it travels the workspace.

B. Torque Control

Manipulators equipped with shaft encoders can often ignore the state-estimation problem when designing a control scheme, as the quality of the state estimate is often independent of the state of the robot. In contrast, the quality of the state estimates inferred from the accelerometers can vary with the configuration and velocity of the manipulator. In this section, we will discuss the ramifications of this property on the behavior of the low-cost manipulator equipped only with accelerometers.

For this experiment, we wrapped a proportional-integral (PI) controller around the state estimates produced by the EKF described in Section III. To reduce the efforts required of the PI controller, we employed active gravity compensation by using the Jacobian to compute the feed-forward torques necessary to float the manipulator links [20].

We found that the accelerometer-only sensing scheme can break down under high dynamic conditions. Specifically, linear accelerations induced by angular joint accelerations are not modeled in the measurement prediction of Equation 3,




Fig. 9. Estimating the repeatability of the manipulator using optical-tracking data of the manipulator repeatedly cycling between two target positions. The plots show the difference between each stopping position of the arm and its respective cluster centroid, in the horizontal plane (left) and a vertical plane (right). The mean deviation distance is 18.6mm. Axes are scaled in millimeters. 14 trials were run, all of which appear on this plot.

and neither are the centripetal accelerations induced by the angular joint velocities. Interestingly, we found that adding those terms did not improve the high-dynamic performance of the manipulator, nor did adding a derivative term to the PI controller. We speculate that for such terms to be effective, the overall system must have a higher bandwidth, calibration level, or measurement SNR.

As a result of the previous observation, the low-cost manipulator could become unstable in high-dynamic situations. Furthermore, the measurement model of Equation 3 does not model contact forces; as a result, large contact forces may also induce instability. In either case, stability can be regained by quickly ramping down motor torques, which effectively slows the manipulator down until it re-enters a stable region of the coupled sensing and control systems.

To quantify the repeatability of the low-cost manipulator, an active optical tracking device (Atracsys EasyTrack500) was used to obtain ground-truth position data of the end effector. The tracking system has an accuracy of 0.2mm. The manipulator was placed in front of the optical tracker, with the optical beacon attached to the gripper tip.

The manipulator was commanded to cycle between two joint configurations which required motion of at least 30 degrees on all six joints of the manipulator (excluding the gripper fingers), resulting in end-effector travel of 34cm. The optical-tracking data was analyzed to extract the points where the manipulator had stopped, resulting in one cluster for each of the target positions. Figure 9 shows an estimate of the repeatability of the manipulator, as measured by the deviation of each of these stopping positions from the mean of its cluster. The mean deviation was 18.6mm, and the maximum deviation was 33.9mm.

C. Position Control

Our final experiment controlled the PR2 Alpha arm in a doorknob-grasping task (Figure 10). To avoid the instabilities witnessed in the low-cost manipulator experiment, we used the accelerometer to control the low-frequency trajectory of the manipulator, and used the PR2's shaft encoders to stabilize the high-frequency behavior. Our accelerometer-based controller sent *relative* joint angle commands to the PR2. As such, this controller could be used without shaft encoders




Fig. 10. In this experiment, accelerometer-based state estimation is used to generate relative joint position commands, allowing a position-controlled robot to repeatedly grasp a doorknob.




Fig. 11. Under relative position control, our control scheme is able to reach out and repeatedly grasp a doorknob. This plot shows a trace of one joint as it undergoes relative joint angle commands from the accelerometer-based sensing scheme and simple setpoint-interpolation to derive small step commands.

on any manipulator with position-based actuators, such as stepper motor-based manipulators. Since the accelerometers fly on the actual links of the robot, their measurements are not corrupted by link droop or backlash.

Our accelerometer-based, encoder-stabilized controller guided the PR2 Alpha through a sequence of control points to repeatedly grasp a door handle in front of the robot. Trajectory tracking in this position-controlled scenario is shown in Figure 11.

VI. DISCUSSION

The work was motivated by the ever-increasing precision of consumer-grade MEMS accelerometers, and the observation that some anticipated future domains of robotics, such as home service robots, possess many sources of error beyond manipulator repeatability. In such scenarios, a reduction in the repeatability of the manipulator may not drastically increase the overall system error.

In general, the accelerometer-based sensing strategy removes complexity from the mechanism and increases the complexity of the calibration and control software. This strategy is motivated by the observation that complex software, unlike complex hardware, can be replicated at no cost.

Because adding accelerometers to an existing manipulator design is mechanically simple and incurs very little cost, this sensing approach is also suitable as a backup, or auxilliary, sensing strategy for manipulators equipped with shaft encoders. The accelerometers could be used to bootstrap the power-up sequence of manipulators: regardless of the configuration of the manipulator at power-up, an accelerometer-driven EKF will quickly converge to a reasonable estimate of the manipulator configuration. After an accelerometer-based joint configuration is estimated, the manipulator could use its incremental encoders to safely and quickly reach the homing flags for each joint. Furthermore, accelerometers can provide information about impacts, drivetrain health (through spectral analysis), and a continual “sanity check” for the incremental encoders.

VII. CONCLUSION AND FUTURE WORK

We presented a low-cost sensing scheme based on 3-d MEMS accelerometers. The system produces coherent, absolute estimates of the kinematic configuration of a manipulator. We performed experiments to quantify the performance of this scheme using both high-precision and low-cost manipulators. The accelerometer-based sensing algorithm can be readily applied to any manipulator to augment its state estimation with very little hardware cost and trivial mechanical complexity.

We intend to continue developing the ideas presented in this paper, exploring a variety of approaches to reduce the cost of actuation, sensing, and control of robotic manipulators in unstructured environments. To gain feedback and encourage replication, the mechanical and electrical designs for the low-cost manipulator, as well as all firmware and software used in the accelerometer-based sensing scheme, will be made freely available at <http://openarms.stanford.edu>

REFERENCES

- [1] A. Saxena, J. Driemeyer, and A. Y. Ng, “Robotic grasping of novel objects using vision,” *International Journal of Robotics Research*, 2008.
- [2] D. Katz, E. Horrell, Y. Yang, B. Burns, T. B. A. Grishkan, V. Zhylkovskyy, O. Brock, and E. Learned-Miller, “The umass mobile manipulator uman: An experimental platform for autonomous mobile manipulation,” *IEEE Workshop on Manipulation for Human Environments*, 2006.
- [3] H. Nguyen, A. Jain, C. Anderson, and C. Kemp, “A clickable world: Behavior selection through pointing and context for mobile manipulation,” *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2008.
- [4] R. Slyper and J. Hodgins, “Action capture with accelerometers,” *Eurographics/ACM SIGGRAPH Symposium on Computer Animation*, 2008.
- [5] E. Foxlin and L. Naimark, “Vis-tracker: A wearable vision-inertial self-tracker,” *IEEE Virtual Reality Conference*, 2003.
- [6] D. Fontaine, D. David, and Y. Caritu, “Sourceless human body motion capture,” *Smart Objects Conference*, 2003.
- [7] L. Bao and S. Intille, “Activity recognition from user-annotated acceleration data,” *Pervasive Computing*, vol. 3001/2004, 2004.
- [8] D. Vail and M. Veloso, “Learning from accelerometer data on a legged robot,” *IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, 2004.
- [9] K. Parsa, J. Angeles, and A. Misra, “Pose-and-twist estimation of a rigid body using accelerometers,” *IEEE International Conference on Robotics and Automation*, 2001.
- [10] F. Ghassemi, S. Tafazoli, P. Lawrence, and K. Hashtrudi-Zaad, “An accelerometer-based joint angle sensor for heavy-duty manipulators,” *IEEE International Conference on Robotics and Automation*, 2002.
- [11] G. Canepa, J. Hollerbach, and A. Boelen, “Kinematic calibration by means of a triaxial accelerometer,” *IEEE International Conference on Robotics and Automation*, 1994.
- [12] H. Aldridge and J.-N. Juang, “Joint position sensor fault tolerance in robot systems using cartesian accelerometers,” *AIAA Guidance, Navigation, and Control Conference*, 1996.
- [13] H. Liu and G. Pang, “Accelerometer for mobile robot positioning,” *IEEE Transactions on Industry Applications*, vol. 37, no. 3, 2001.
- [14] Y. F. Li and X. B. Chen, “End-point sensing and state observation of a flexible-link robot,” *IEEE/ASME Transactions on Mechatronics*, vol. 6, no. 3, 2001.
- [15] N. Miller, O. Jenkins, M. Kallman, and M. Mataric, “Motion capture from inertial sensing for untethered humanoid teleoperation,” *International Journal of Humanoid Robotics*, 2008.
- [16] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [17] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, 2000.
- [18] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Y. Ng, “Ros: an open-source robot operating system,” *Open-Source Software workshop of the IEEE International Conference on Robotics and Automation*, 2009.
- [19] K. Wyrobek, E. Berger, H. V. der Loos, and J. K. Salisbury, “Towards a personal robotics development platform: Rationale and design of an intrinsically safe personal robot,” *IEEE International Conference on Robotics and Automation*, 2008.
- [20] J. Craig, *Introduction to Robotics: Mechanics and Control*, 3rd ed. Pearson Prentice Hall, 2005.