Natural Task Decomposition with Intrinsic Potential Fields

Stephen Hart and Roderic Grupen
Department of Computer Science
University of Massachusetts Amherst
Amherst, MA 01003 USA
{shart | grupen}@cs.umass.edu

***Abstract*— Any given task can be solved in a number of ways, whether through path-planning, modeling, or control techniques. In this paper, we present a methodology for natural task decomposition through the use of *intrinsically* meaningful potential fields. Specifically, we demonstrate that using classical conditioning measures in a concurrent control framework provides a domain-general means for solving tasks. Among the conditioning measures we use are manipulability [1], localizability [2], and range of motion. To illustrate the value of our approach we demonstrate its applicability to an industrially relevent inspection task.**

I. Introduction

Nikolai Bernstein defined *dexterity* in terms of an organism's versatility—both physically and cognitively—to adapt to a rich and unpredictable environment [3]. In this paper, we demonstrate the usefulness of conditioning kinematic chains to respond competently under such uncertainty. As a result, we show how traditional conditioning measures can be used in a non-traditional way: as plans to solve behavioral problems. Such techniques are not specific to a task, but instead provide *intrinsic* fields that can be applied in manipulation and recognition tasks.

One such plan for posturing a system is through *isotropic* conditioning which allows for a tradeoff in the ability to perceive (input) errors and to impart (output) movements. In such a way, isotropic conditioning exhibits a principle of flexibility and least commitment for unknown future circumstances. Such a methodology is therefore useful when programming dexterous robots that must achieve a variety of complex behavior in uncertain real-world environments. We also present kinematic *constraints*—such as keeping a manipulator away from its joint range limits—as conditioning fields. All conditioning fields presented in this paper, however, are similar in that they provide a system with a potential function that directs the mechanism toward desirable kinematic configurations.

In this work, conditioning fields are used directly as potential fields which conditioning controllers can follow toward "sweet-spots" where the device can yield optimal kinematic properties. When employed in a concurrent control framework, subordinate conditioning controllers may increase the efficiency of higher-priority, task-specific, position or force controllers. For example, in the presence of uncertainty, it is beneficial to keep a manipulator kinematically isotropic when moving to grasp an object of unknown geometry. When conditioning controllers are used as superior objectives, such as moving an acquired object to a well-conditioned visual sweet-spot, they provide goals that may increase the effectiveness and accuracy of a task, such as feature recognition. Conditioning controllers are formulated in this paper using the *control basis* framework for multi-objective control.

II. Related Work

The condition number of the manipulator Jacobian was introduced by Salisbury and Craig [4] to evaluate the kinodynamic state of a mechanism. Yoshikawa extended these ideas to provide a means of optimizing the *manipulability* of a mechanism while performing a task [1]. The proposed method biases the state of the manipulator to more isotropic configurations where forces and velocities can be applied equally in any direction. Examples of mechanisms employing kinematic and isotropic conditioning techniques based on the manipulator Jacobian can be found in [5], [6], [7], [8], [9]. The concept of isotropic conditioning is equally effective for any linear transformation and was later generalized to acceleration and inertial measures [10], [11], [5], [12], and also to evaluate viewpoint quality in a stereo system in order to maximize localization precision [2]. In this paper, we demonstrate how kinematic conditioning underlies both the manipulability and localizability concepts in the previous literature and the role of both of these concepts in dexterous inspection tasks.

Using the same techniques, anisotropic conditioning of a mechanism can be employed when the task is well-specified ahead of time. Chiu demonstrated how to utilize the redundancy of a manipulator to posture an end-effector to optimize the application of velocities and forces along known task directions [13]. In this work, it was shown how to increase either amplification or precision along a specified task-direction. In this paper, we assume little prior task knowledge and focus on isotropic conditioning only.

Section IV provides a description of the control basis framework for concurrent control [14], [15]. This framework has been used in many domains including grasp-control, multi-robot mapping and localization, and walking gait formulation [14], [16], [17]. Section V presents a means of representing conditioning metrics in terms of the control basis.

III. Isotropic Conditioning

The "condition" of a linear transformation, $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}$, can be described in terms of the deformation $\mathbf{A}$ causes to input signals, $\boldsymbol{x}$. The influence of transformation $\mathbf{A}$ can

be visualized by mapping an m-dimensional hyper-sphere "test pattern" $\|\boldsymbol{x}\|^2 = x_0^2 + x_1^2 + \ldots + x_m^2 \leq 1$ through an $n \times m$ transform $\mathbf{A}$ to see how it is distorted into an n-dimensional *hyper-ellipsoid*. This "conditioning" ellipsoid reveals the directional dependence of the transformation, $\mathbf{A}$ [5]. Figure 1 portrays such an analysis in two dimensions.

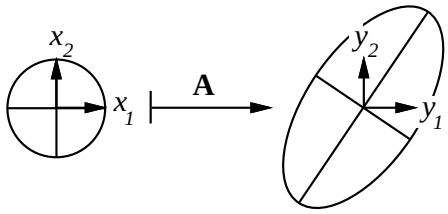

Fig. 1. The distribution of inputs defined by the circle in the $\mathbf{x}$ plane, $\boldsymbol{x}^T\boldsymbol{x} \leq 1$, maps through $\mathbf{A}$ to an ellipse in the $\boldsymbol{y}$ plane defined by the singular values and singular vectors of $\mathbf{A}$, or equivalently, the eigenvalues and eigenvectors of $\mathbf{A}\mathbf{A}^T$.

If output movements are unknown for a particular task *a priori*, isotropic conditioning can improve the ability of the manipulator to respond uniformly to unpredictable events in unknown geometries. Isotropic conditioning can be accomplished by examining the volume of the conditioning ellipsoid. The volume of the ellipsoid can be appreciable even when the matrix $\mathbf{A}$ is near a singularity, but in general, volume increases as the conditioning ellipsoid becomes more spherical.

In order to derive an analytical description of the conditioning ellipsoid's volume, note that the singular vectors of $\mathbf{A}$ are the eigenvectors of $\mathbf{A}\mathbf{A}^T$ and the singular values of $\mathbf{A}$ are the square roots of the eigenvalues of $\mathbf{A}\mathbf{A}^T$. The ellipsoid is expressed in terms of its singular values $\Sigma = [\sigma_1 \ \sigma_2 \ \cdots \ \sigma_m]$, in descending order of magnitude. The product of the singular values, $\prod_{i=1}^{m} \sigma_i$, yields a measure proportional to the volume of the ellipsoid. If $\mathbf{A}$ is singular, its volume is zero. Alternatively, the same measure can be computed

$$\kappa(A) \quad = \quad \sqrt{det(\mathbf{A}\mathbf{A}^T)}. \tag{1}$$

A. Measure of Manipulability

Yoshikawa applied the above analysis to the manipulator Jacobian $\mathbf{J_m}$ and proposed a *measure of manipulability* (MoM) that evaluates the kinematic "condition" of a robot mechanism. Effectively capturing the distance the mechanism is from a singularity, this measure can keep the system well-conditioned to impart velocities or forces in any direction. By Equation 1, the measure of isotropy of the manipulator Jacobian is defined as:

$$\kappa_m \quad = \quad \sqrt{det(\mathbf{J}_m\mathbf{J}_m^T)}. \tag{2}$$

B. Measure of Localizability

We can also pose the stereo reconstruction of a unique feature on two image planes as a kinematics problem and apply the above conditioning analysis. The stereo triangulation equations are used to transform an oculomotor configuration into a Cartesian coordinate of the feature. We can compute the principal kinematic transformations to evaluate visual acuity as a function of the relative position of the feature. We can construct an oculomotor Jacobian for the stereo triangulation equations, $d\boldsymbol{r} = \mathbf{J}_v d\boldsymbol{\gamma}$, the transformation from oculomotor to Cartesian velocities. As such, the oculomotor Jacobian may be used to estimate the *sensitivity* to errors in the observed configuration, $[\gamma_L \ \gamma_R]$. Applying Equation 1 to the oculomotor Jacobian $\mathbf{J}_v$ we can compute the visual condition of a stereo configuration:

$$\kappa_l \quad = \quad \sqrt{det(\mathbf{J}_v\mathbf{J}_v^T)} \tag{3}$$

The oculomotor conditioning ellipsoid represents the error covariance in stereo localization and consequently, the spatially anisotropic uncertainty of the stereo imaging geometry. We will thus define the *measure of localizability* MoL as κ_l. Figure 2 shows the MoL for a stereo camera pair in 2-dimensions. This scalar field is white where the stereo geometry is well conditioned for precise feature localization.

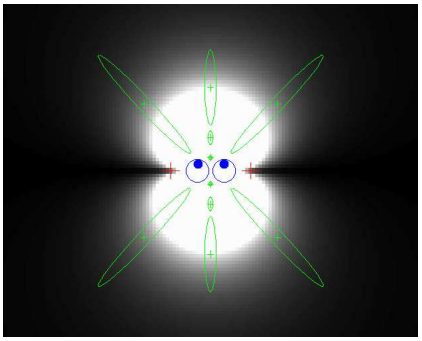

Fig. 2. The Relative Localizability Scalar Field: for the 2-D case $det(\mathbf{J}_v\mathbf{J}_v^T)^{1/2}$. Examples of the localizability ellipsoid at several positions marked by crosshairs are overlaid onto the field.

Superimposed on the localizability field in Figure 2 are ellipsoids derived from $\mathbf{J}_v\mathbf{J}_v^T$ at several positions marked by crosshairs. The ellipsoids illustrate the shape (if not the magnitude) of the Cartesian error covariance. Generally, the lateral error is relatively small and the radial error can vary dramatically depending on the position of the feature.

IV. The Control Basis

The conditioning fields presented above will be used directly as artificial potentials. Control inputs in the quasistatic case can be computed as the gradient of these conditioning potentials. We will formulate these controllers using the control basis framework for discrete event dynamic systems [14], [15], [18]. Controller i is denoted ϕ_i.

Concurrent control commands are constructed by projecting the output of a subordinate controller, ϕ_2, into the nullspace of a higher priority controller, ϕ_1 The nullspace $\mathcal{N}_1$ of the control command of ϕ_1 is $(I - J_1^{\#}J_1)$ where J_i is the Jacobian matrix of the objective with respect to the configuration variables θ and $J_i^{\#}$ is its pseudoinverse [5].

Nullspace projections can be chained together to capture an array of prioritized control objectives. Our shorthand for this nullspace projection is written using the "subject-to" operator "$\triangleleft$" [14]. For example, $\phi_2 \triangleleft \phi_1$—read, "$\phi_2$ *subject-to* ϕ_1"—captures the case where subordinate controller ϕ_2 projects outputs into the nullspace of the superior controller ϕ_1.

V. Controller Definitions

In this section, we define controllers that allow for a mechanism to optimize the system's configuration according to the conditioning metrics described above.

A. Isotropic Conditioning Controllers

For an isotropic manipulator conditioning field, we define a potential field according to its gradient:

$$p_{kc}(\mathbf{q}) = -\frac{d}{d\mathbf{q}}\sqrt{det(\mathbf{J}_m\mathbf{J}_m^T)} \tag{4}$$

We define a kinematic conditioning controller ϕ_{kc} that computes an error according to the gradient of this field, and moves to maximize the system's isotropic conditioning.

For the oculomotor Jacobian $\mathbf{J}_v(\gamma)$, we can describe a potential field according to the gradient of the corresponding *relative* localizability scalar field:

$$p_{vc}(\gamma) = -\frac{d}{d\gamma}\sqrt{det(\mathbf{J}_v\mathbf{J}_v^T)} \tag{5}$$

We define a visually conditioning controller ϕ_{vc} that computes an error according to this gradient, and moves to maximize the system's isotropic visual acuity. The ϕ_{vc} controller can be "dexterous"—it can engage either manipulator or head degrees of freedom to ascend the field.

B. Constraint Posturing Controllers

In addition to conditioning objectives that follow gradients defined by the *shape* of the Jacobian transformations, we can apply other objective fields that posture the mechanism towards desirable kinematic configurations. In particular, we will describe a metric that measures proximity to joint range limits and a metric that measures proximity to self-collisions.

For an n-dimensional mechanism, we can define a n-dimensional cosine field around the center of each joint's range of motion. The field for a set of joint angles $\mathbf{q}$ is defined by:

$$c_{rom}(\mathbf{q}) = -\sum_i^n cos(r_i(q_i)) \tag{6}$$

where

$$r_i(q_i) = \frac{q_i - \bar{q_i}}{q_{max} - q_{min}}\pi. \tag{7}$$

In this equation, q_{max} and q_{min} represent the upper and lower limits of q_i's range of motion, and $\bar{q_i} = (q_{max} - q_{min})/2$. This field provides a convex potential that is centered halfway between each joint's upper and lower limits. We can define a controller ϕ_{rom} that moves a set of joints according to the gradient of this field:

$$p_{rom}(\mathbf{q}) \quad = \quad -\frac{d}{d\mathbf{q}}c_{rom}(\mathbf{q}) \tag{8}$$

It may also be useful to define an objective function that prevents a mechanism from self-collisions. Ideally, this function f_{obs} may be harmonic or sub-harmonic. In the experiments presented in this paper we use a simple potential field approach [19] and define the controller ϕ_{obs}.

C. REACH Controller Composition

In this section, we describe three REACH composite controllers that allow a robot to optimize the conditioning metrics described above in the presence of other task objectives.

1) Reaching to a Target Location: Often, we desire a manipulator to move to a desired location in Cartesian space. A control law can be defined using operational space motion control to bring a system to a given reference [20]. We will define such a controller as ϕ_{pos}. We optimize the metrics described above as much as possible by projecting the conditioning objectives into the *nullspace* of ϕ_{pos}. We define the first REACH controller as:

$$\phi_{R1} = \phi_{rom} \triangleleft \phi_{kc} \triangleleft \phi_{pos} \triangleleft \phi_{obs} \tag{9}$$

2) Multi-Objective Conditioning Controllers: Furthermore, we can describe *intrinsic* goals of the system that bring the manipulator to well-conditioned locations. The following two control laws bring the system to kinematic and oculomotor "sweet-spots" while still preventing collisions.

$$\phi_{R2} = \phi_{rom} \triangleleft \phi_{kc} \triangleleft \phi_{obs} \tag{10}$$

and

$$\phi_{R3} = \phi_{rom} \triangleleft \phi_{kc} \triangleleft \phi_{vc} \triangleleft \phi_{obs} \tag{11}$$

These three control laws provide useful actions that can be assembled into policies based on task requirements. It is important to note the prevalence of conditioning metrics in these policies which we will see examples of in the next section.

VI. Experiments

We now present demonstrations in which isotropic and constraint conditioning controllers are assembled in a concurrent control framework. The demonstrations were performed on the the bimanual humanoid robot "Dexter," seen in Figure 3. This robot has a 2-DOF pan/tilt head equipped with two Sony cameras capable of stereo triangulation and two 7-DOF Whole-Arm Manipulators (Barrett Technologies, Cambridge MA). Each WAM is equipped with a 3-finger, 4 DOF Barrett Hand and 6-axis finger-tip load cells.

In the first demonstration, manipulability is optimized in the course of reaching to an object of unknown shape such that a grasp controller may achieve wrench closure. Next, a concurrent control law is used to condition the state of the mechanism subject to maintaining a grasp on the object. Lastly, it is shown how the localizability metric may be used to inform action selection depending on the character of the visual features to be inspected. Salient features of some

objects are inspectable immediately, others must be grasped and moved to condition the stereo vision system further.

Concurrent control laws described in Section V-C, are used to reach to well-conditioned locations either exclusively or subject to reaching to a target location. The kinematic constraint controllers are also employed to ensure collision-free movements and limitations with respect to joint ranges of motion. The result is that conditioning control is subordinate to every control task submitted to the robot.

A. Grasping an Object

Previous work by Coelho [18] and Platt [17] provided a control formulation for achieving a wrench closure grasp on an object with shape unknown *a priori*. The resulting controller, ϕ_{wc} minimizes the net wrench residual between the fingers in contact with the object. This controller requires the dexterous application of movements and forces, and demonstrates the necessity for conditioning the mechanism before and during grasp controller execution when these movements and forces are not known ahead of time.

In our first set of experiments, the robot begins in a start configuration determined by the execution of ϕ_{rom}. This position is convenient because it places both arms above the table as seen in Figure 3(a), but does not obstruct the stereo head's view of any object placed on the table. The robot then reaches both of its arms towards a box placed in its reachable workspace. The position of the box is recovered through stereo triangulation and background subtraction techniques. When localization completes, the robot employs control law ϕ_{R1} to bring its arms close to the box, as seen in Figure 3(b). At this point, ϕ_{wc} takes over to achieve wrench closure.

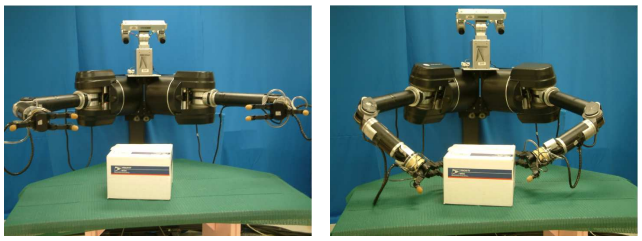

(a) (b)

Fig. 3. (a) and (b) show the robot after localizing the box before and after, respectively, having completed a bimanual reach to its location.

B. Conditioning the Grasp

Once wrench closure is achieved, the arm configuration is optimized for the average manipulability of both arms subject to maintaining the grasp. The initial grasp is governed by compliance to the object geometry. However, after a grasp is achieved, the robot may reconfigure to better, kinematically isotropic, arm postures while preserving the wrench closure.

Figures 4(a) and 4(b) show the robot before and after the execution of the new control combination:

$$\phi_{cg} \quad = \quad \phi_{R2} \triangleleft \phi_{wc} \tag{12}$$

Figures 4(c) and 4(d) show the trajectory of the robot's arms and the box for two instances of this control law. Figure 5 shows the increase of the MoM averaged over six executions of ϕ_{cg} performed on the robot after grasping the box from different starting positions and orientations, all within the robot's reachable workspace. Shown in the graph is the MoM for each arm individually, as well as their average. The average MoM was the metric optimized, as can be seen by the signal's monotonic increase.

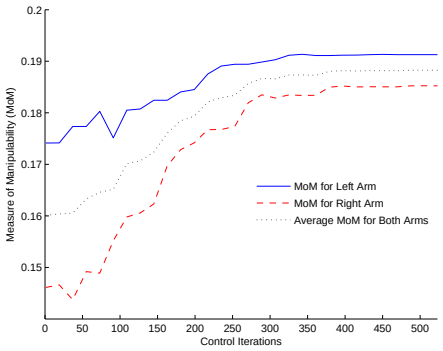

Fig. 5. This graph shows the increase in manipulability that occurs through the execution of ϕ_{cg} after the box has been grasped. The plot shows the measure of manipulability (MoM) for each arm individually, and averaged over 6 trials performed on the robot after grasping the box from different starting positions and orientations.

C. Object Identification

In the final example of the usefulness of kinematic conditioning, we demonstrate a controller designed to move the grasped object into a visually conditioned "sweet-spot" such that a barcode pattern on the object's surface can be inspected and classified. This task requires a visual perspective compatible with the required spatial precisions. The robot can achieve stereo configurations that provide the appropriate visual acuity by ascending the localizability field. Figure 6 shows how the measure of localizability metric κ_l increases according to control law ϕ_{R3} for both one- and two-handed grasps. Figure 7 shows the robot after the execution of this control law while maintaining a two-handed grasp.

For these experiments, a barcode pattern belonging to one of three different sets was placed on the robot-facing side of the box used in the previous demonstrations. Each set of patterns contains three test-patterns that have the same maximum spatial frequency bandwidth, designated as either "high," "low," or "medium." Figure 9 shows the three sets of three patterns. Table I shows the width of the smallest period T of each barcode pattern, as well as their expected pixel resolution on the robot's cameras and the MoL metric at three different depths x_d. These values were calculated for the 640x480 images using a pinhole camera model with a focal length of 837 mm. The depths correspond to the location of the box 1) in the center of the table in front of the robot, after localizing it, 2) after the robot has grasped the box using control law $\phi_{R2} \triangleleft \phi_{wc}$ (using a two handed grasp), and 3) after the execution of control law $\phi_{R3} \triangleleft \phi_{wc}$. The views from the robot's left camera at each stage of the policy are seen in Figure 8.

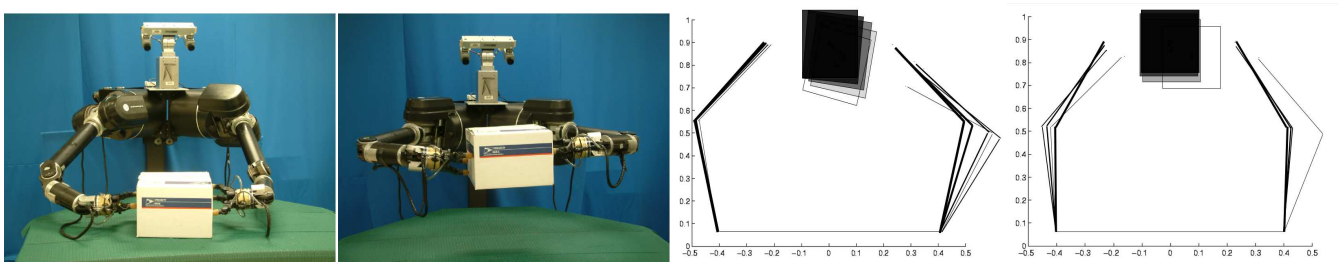

(a) (b) (c) (d)

Fig. 4. (a) and (b) show pictures of the robot before and after the execution of ϕ_{cg}. (c) and (d) show top-down representations of two sample trajectories the robot follows during execution of the same control law. Note that the confi gurations of the robot's hands (joining the arms to the object) are not shown. These graphs were generated from actual position information. Time progresses from light to dark.

(a) (b) (c)

Fig. 8. Screen-shots from the robot's left camera (a) with the box on the table, (b) with the box grasped with control law $\phi_{R2} \triangleleft \phi_{wc}$, and 3) following the execution of $\phi_{R2} \triangleleft \phi_{wc}$. These 640x480 pixel images were captured while one of the high-frequency patterns was placed on the box.

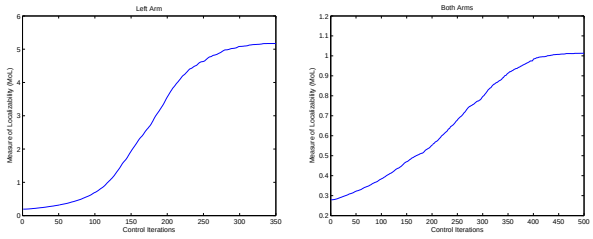

(a) (b)

Fig. 6. (a) shows the increase of the measure localizability metric under ϕ_{R3} while grasping a small object. (b) shows the same metric when the robot grasps the larger box with two hands. Notice how the increase is greater for the smaller object. This performance gain is a result of being unconstrained by the wrench closure controller, which in the two handed case has to maintain the relative positions and orientations of both hands.

TABLE I

PIXELS PER PERIOD AND MOL

	x_d = 100 cm	x_d = 75 cm	x_d = 50 cm
T_{high} = 2 mm	1.67	2.32	3.35
T_{medium} = 6 mm	5.02	6.97	10.04
T_{low} = 30 mm	25.11	33.48	50.02
MoL	1.05	2.32	7.31

For each pattern in each test set and for the pattern perceived from the robot's left camera, the barcode was converted to a string of intervals based on color and bar width. Classification was then performed by finding the lowest mean-squared error match between the perceived test pattern and the three patterns in the same (known) frequency class. Figure 10 shows the accuracy in classification for each pattern set at each of the three locations (on the table, at the most bimanual manipulable position, and where there is the most visual acuity, subject to self-collisions). Each classification bar is the average result from ten trials. It can be seen that the ability to perceive higher frequencies clearly increases when moving the object towards the visually conditioned "sweet-spot." Note that, due to the symmetry of the bimanual system, the execution of control law $\phi_{R2} \triangleleft \phi_{wc}$ also brings the object to a more visually isotropic configuration then when the object is on the table. These experiments suggest that, if the spatial frequency of the test patterns is known, the robot can actively decide which action needs to be taken—how close the robot needs to bring the object to its cameras—to accurately perform classification.

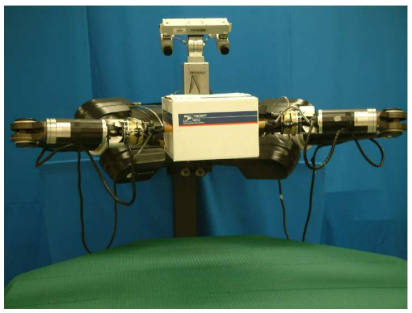

Fig. 7. This picture shows the robot after the execution of control law $\phi_{R3} \triangleleft \phi_{wc}$.

VII. DISCUSSION AND CONCLUSIONS

In this paper we introduced a uniform treatment of conditioning of the manipulator and oculomotor Jacobians. The

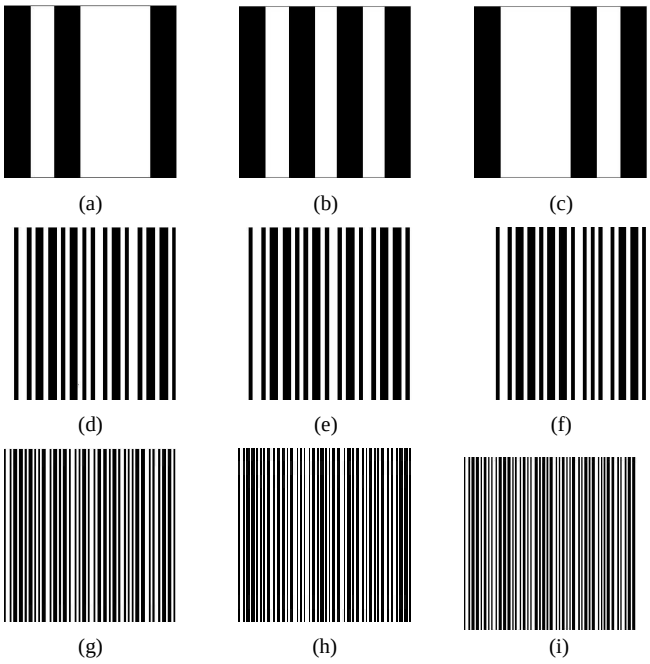

Fig. 9. The barcode patterns used for classification experiments. Patterns (a-c) are the low-frequency patterns, and were randomly generated. (d-f) are medium-frequency patterns, representing the characters A, B, and C, respectively, in the Code-39 barcode standard. The high-frequency patterns (g-i), are the Code-39 encoded strings ROBOT, DEXTER and AMHERST, respectively. Each pattern is 10 cm square.

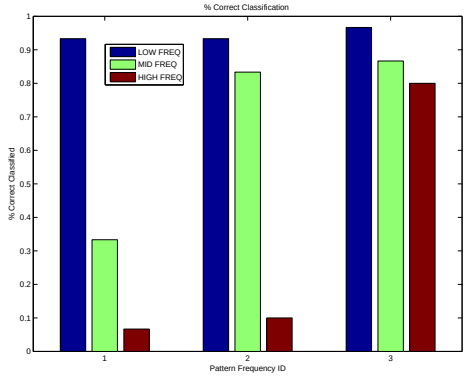

Fig. 10. The classification error of each of the barcode pattern groups while the box is on the table, at the most bimanually manipulable position, and at the most visually acute position. It is clear that classification of the low-frequency box is successful while the object is on the table. After moving the object closer, using $\phi_{R2} \triangleleft \phi_{wc}$, the low- and medium-frequency patterns are classifiable. After the execution of $\phi_{R3} \triangleleft \phi_{wc}$, where visual acuity is maximal, all three frequency patterns are classifiable with a high degree of accuracy.

resulting metrics provide objectives that can be attached as subordinate goals to other behavior, or as that can be used as superior objectives to bring the system to intrinsic "sweet-spots." Furthermore, we demonstrate how controllers designed to optimize these metrics, along with task-specific, and constraint-driven controllers, can be employed in a combinatoric basis to achieve a wealth of "natural" behavior for a robot system. The control laws provide important behavior that a robot operating in an unknown environment may employ in light of uncertain future circumstances.

The ideas presented in this paper also suggest a method for performing sensor fusion to accurately localize the position or size of grasped objects. Computing the conditioning ellipsoids of visual and spatial acuity metrics provide the first and second moments needed for the application of Kalman Filter estimation techniques. Although, a demonstration of this idea is beyond the scope of the current paper, it is the subject of current research.

ACKNOWLEDGMENTS

This research was supported under contract numbers ARO W911NF-05-1-0396 and NASA NNJ05HB61A-5710001842 and under NASA Graduate Student Research Program fellowship NNJ05JG73H.

REFERENCES

[1] T. Yoshikawa, "Manipulability of robotic mechanisms," *The International Journal of Robotics Research*, vol. 4, pp. 3–9, 1985.

[2] J. Uppala, D. Karuppiah, M. Brewer, C. Ravela, and R. Grupen, "On viewpoint control," in *International IEEE Conference on Robotics and Automation*, 2002.

[3] N. Bernstein, *Dexterity and its Development*, M. Latash and M. Turvey, Eds. Mahwah, N.J.: Laurence Erlbaum Associates Inc., 1996.

[4] J. Salisbury and J. Craig, "Articulated hands: Force control and kinematic issues," *The International Journal of Robotics Research*, vol. 1, pp. 4–17, 1982.

[5] Y. Nakamura, *Advanced Robotics: Redundancy and Optimization*. Addison-Wesley, 1991.

[6] H. Asada and J. C. Granito, "Kinematic and static characterization of wrist joints and their optimal design," in *International IEEE Conference on Robotics and Automation*, 1985.

[7] J. Angeles and C. Lopez-Cajun, "Kinematic isotropy and conditioning index of serial robotic manipulators," *The International Journal of Robotics Research*, vol. 11, pp. 560–571, 1992.

[8] R. Grupen and K. Souccar, "Manipulability-based spatial isotropy: A kinematic reflex," in *Proceedings of the International Workshop on Mechatronical Computer Systems for Perception and Action*, 1993, pp. 157–163.

[9] F. Ranjbaran, J. Angeles, and A. Kecskemethy, "On the kinematic conditioning of robotic manipulators," in *International IEEE Conference on Robotics and Automation*, Minneapolis, MN, 1996.

[10] P. Chiacchio, S.Chiaverini, L.Sciavicco, and B.Siciliano, "Influence of gravity on the manipulability ellipsoid for robot arms," *Journal of Dynamic Systems, Measurement, and Control*, vol. 114, no. 4, pp. 723–727, 1992.

[11] P. Chiacchio, "A new dynamic manipulability ellipsoid for redundant manipulators," *Robotica*, vol. 18, no. 4, pp. 381–387, 2001.

[12] M. Rosenstein and R. Grupen, "Velocity-dependent dynamic manipulability," in *International IEEE Conference on Robotics and Automation*, Washington, DC, 2002.

[13] S. Chiu, "Control of redundant manipulators for task compatability," in *International IEEE Conference on Robotics and Automation*, 1987.

[14] M. Huber, "A hybrid architecture for adaptive robot control," Ph.D. dissertation, Department of Computer Science, University of Massachusetts Amherst, 2000.

[15] S. Hart, S. Ou, J. Sweeney, and R. Grupen, "A framework for learning declarative structure," in *Workshop on Manipulation for Human Environments, Robotics: Science and Systems*, Philadelphia, PA, July 2006.

[16] J.Sweeney, T. Brunette, Y.Yang, and R. Grupen, "Coordinated teams of reactive mobile platforms," in *International IEEE Conference on Robotics and Automation*, 2002.

[17] R. Platt, "Learning and generalizing control based grasping and manipulation skills," Ph.D. dissertation, Department of Computer Science, University of Massachusetts Amherst, 2006.

[18] J. A. Coelho, "Multifingered grasping: Grasp reflexes and control context," Ph.D. dissertation, Department of Computer Science, University of Massachusetts Amherst, 2001.

[19] J. Latombe, *Robot Motion Planning*. Norwell, MA: Kluwer Academic Publishers, 1991.

[20] J. Craig, *Introduction to Robotics: Mechanics and Control*, 3rd ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2005.