Manipulation Manifolds: Explorations into Uncovering Manifolds in Sensory-motor Spaces

Odest Chadwicke Jenkins
Computer Science Department
Brown University
Provicence, RI, USA 02910-1910
cjenkins@cs.brown.edu

Bobby Bodenheimer, Richard Alan Peters II
Center for Intelligent Systems
Vanderbilt University School of Engineering
Nashville, TN 37235
bobbyb@vuse.vanderbilt.edu, Alan.Peters@Vanderbilt.Edu

***Abstract*— This paper presents results from the application of dimensionality reduction algorithms to sensory-data time-series that were recorded from Robonaut – NASA's humanoid robot – while it was being teleoperated through four tool manipulation tasks. The algorithms tested were Principal Component Analysis, Multidimensional Scaling, and Spatio-Temporal Isomap. Structures were shown to exist in some cases, but their detection required careful analysis and a correct choice of parameters.**

***Index Terms*— Spatio-temporal learning, dimension reduction, robotics, dexterous manipulation.**

I. Introduction

As observed by Pfeifer, sensorimoter coordination (SMC) data can self-organize into descriptors that categorize the interaction of a robot with its environment [17]. Similarly, our focus is on the mapping from sensory observations to motor commands, the essence of robot control. This mapping defines a structure in the robot's sensorimotor state space (SMSS). Understanding and visualizing such a structure would potentially be of great value in improving control policies for a complex robot. Simply watching a robot's sensorimotor data stream during operation often reveals patterns that are indicative of a structure. Unfortunately, the robot's sensorimotor state space is high-dimensional, having dimension equal to the number of scalar observations (sensory and motor) relevant to the control of the robot, which makes visualizing and understanding structures in the SMSS difficult. However, it is likely that the intrinsic dimensionality of control and action strategies in the SMSS is of a much lower dimensionality, and in fact, describes a manifold immersed in the SMSS. In this paper, our goal is to employ dimensionality reduction and manifold learning techniques to understand the manifold structure of control and action strategies for a robot.

Many robot tasks (such as dexterous manipulation) may have latent sensorimotor structures that could be uncovered through manifold learning. In repetitive, constrained motion by the robot (e.g., repeatedly reaching toward and grasping an object), the dominant vari- ables tend to trace low-dimensional nonlinear manifolds in the SMSS. Repetitive variations of a task under different conditions lead to a family of trajactories that lie on a manifold surface of the SMSS. Given the sensorimotor data from these variations in the SMSS, in this paper we apply several dimensionality reduction techniques to learn the structure of the lower-dimensional manifolds. In particular, we employ PCA [10], MDS [9], Isomap [21], and Spatio-Temporal Isomap [12] to learn the manifold strucutre. Each of these techniques learns the manifold differently, and thus each may reveal different aspects of the high dimensional structure. Determining the strengths and weaknesses of these algorithms as applied to robotic data is a further goal of this work. In particular, we present results from the application of these techniques to sensorimotor time-series recorded during the teleoperation of the NASA Robonaut. Experimental results are provided for four sets of trials pertainting to teleoperated manipulation.

As mentioned above, knowledge of the SMSS manifolds that correspond to various robotics tasks and scenarios would be valuable. We focus specifically on grasping and manipulation tasks for humanoid robots. Uncovering manifolds for such tasks may lead to compact descriptions of fundamental behaviors (encapsulation or macro-generation), real-time behavior interpolation, recognition of key events or unexpected events during task execution, recognition of intentionality in the perceived behavior of a collaborator (human or robot), or the generalization of related tasks through the detection of vector space homeomorphisms. Furthermore, many of these applications could be cast in terms of probabilistic reasoning, such as particle filtering [23], where uncovered sensorimotor structures as priors on sensor-motor experiences. In combination with clustering, functional or numerical descriptions of these manifolds could also uncover sensory-motor categories and lead to learning algorithms for sensorimotor coordination and goal attainment.

II. Related Work

In [16] a single SMSS trajectory was learned over six trials that could later be performed autonomously with success in the face of small variations in the environment or perturbations of the goal. Later, it was shown that sets of such learned trajectories could be interpolated to provide intermediate

results [6]. In addition to Pfeifer [17], many others have studied the extraction of SMC parameters, including Cohen [8], Grupen [7], Lungarella [13], and Peters [5].

The use of motion data to plan robotic motion is a problem that has been studied by various groups, such as efforts by Matarić [15], Jenkins and Matarić [11], Ude et al. [24], Pollard et al. [18], and Atkeson et al. [3]. Similar to this work, Matarić and Jenkins have used Spatio-temporal Isomap towards the separate topic of learning predictive motion primitives from motion capture for controlling simulated humanoids. Isomap [22] is one of a number of dimensionality reduction techniques including Principal Component Analysis and the related Singular Value Decomposition, both of which are textbook approaches, Multi-dimensional Scaling [9] and Locally Linear Embedding [20]. Application of dimensionality reduction to robotics has been studied by many, including Asada [2], MacDorman [14], and others.

An alternative, but complimentary, approach to uncovering sensory-motor structure is to treat is as a regression between input (sensory) and output (motor) spaces. Several papers by Atkeson and Schaal, including [4] address this regression problem through the use of receptive fields. Discovery of manifold structures could help guide the placement or repositioning of such receptors given their sensitivity to order due to incremental inference.

III. Background in Dimension Reduction

We assume sensory-motor observables, given as vectors at ordered discrete instances of time $\mathbf{x}_k : \{1, 2, 3, \ldots N\} \rightarrow \mathbb{R}^D$, are intrinsically parameterizated by a lower dimensional embedding. The embedding provides a mapping $\mathbf{x} = \phi(\mathbf{y})$ between intrinsic parameters and observations, realizing intrinsic coordinates $\mathbf{y}_k : \{1, 2, 3, \ldots N\} \rightarrow \mathbb{R}^d$ for the input data where $d < D$. Arguably the most popular and well-known dimension reduction technique is Principal Components Analysis (PCA) [10]. PCA involves an eigendecomposition on a linear covariance matrix to find an orthogonal subspace of principal components compactly approximating the input data.

Multidimensional scaling (MDS) [9] is another approach where pairwise distances, rather than linear covariance, are preserved. Given the distance between all input datapairs $D_{\mathbf{x}_i, \mathbf{x}_j}$, MDS produces embedding coordinates that minimizes the error $E = |D_x - D_y|_{L^2}$, where D_x and D_y are respectively the pairwise input and embedding space distance matrices. Essentially, MDS produces embedding coordinates that preserves thedistance metric as much as possible. MDS is equivalent to PCA when the input space distance metric is Euclidean[25]. Isomap [22] uses a geodesic (Dijkstra shortest-path) distance metric with MDS, avoiding "short-circuiting" problems.

These techniques, however, are not necessarily suited for time-series data because they assume the input are i.i.d., independent samples from the same manifold parameterization.

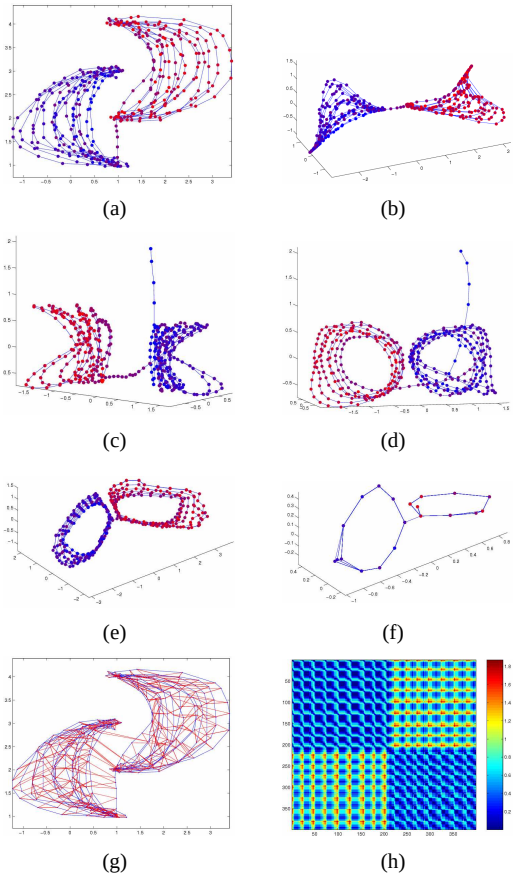

Fig. 1. (a) A 2D mouse-click trajectory color coded from blue to red to indicate progress through time. (b) Embedding of the trajectory using standard (non-ST) Isomap. (c,d) Two views of the trajectory embedded by temporally-windowed MDS. (e,f) Two ST-Isomap embeddings of the trajectory showing a collapse of the trajectories that results from higher c_{ctn} values. (g,h) The spatio-temporal correspondences and the spatio-temporal distance matrix computed by ST-Isomap.

Time-series data are not independent, but rather sequentially ordered samples from an underlying spatio-temporal process. One means to account for this temporality is to add time as another dimension to the input data. We explore the use of a "windowed MDS" procedure, where each input data object is a temporally extended window of observations. Adding time as another dimension serves to disambiguate spatially proximal datapairs indicative of different phases of a temporal process, but does not account for corresponding spatially distal datapairs representative of equivalent phases in the underlying temporal process. As we describe in next subsection, spatio-temporal Isomap (ST-Isomap) [12] accounts for both proximal disambiguation and distal correspondence (cf. section III-A).

A simple experiment illustrates the differences in these approaches (Fig. 1). As input, two 2D "moons" (crescent

shapes) were formed from a sequence of mouse clicks repeatedly traversing the pattern. PCA performs an affine transformation of the data. Isomap disregards the temporal sequencing and pushes the crescents away from each other. Windowed MDS disambiguates each crescent traversal into a loop. ST-Isomap uncovers a more compact description of each loop by registering spatio-temporal correspondences.

A. Uncovering Manifolds with Spatio-temporal Isomap

A spatio-temporal manifold, an example of which is shown in Fig. 1, can be described structurally as a $1+\epsilon$-manifold curve in an embedding space. Successive locations on this manifold represent successive time-instants in the spatio-temporal process. Each location on this manifold encapsulates all of the spatial variations representing a certain phase along the spatio-temporal process. Spatial variations being different locations in input space that correspond to a particular instance of space-time along the manifold. Thus, a diverse set of spatial variations that correspond to the same phase is collapsed into a single location in the embedded manifold.

The general procedure for ST-Isomap is:

1) compute a sparse L_2 distance matrix D^l from local neighborhoods $\texttt{nbhd}(x_i)$ about each point x_i;
2) locally identify the *common temporal neighbors* $\texttt{CTN}(x_i)$ of each point x_i as the *K-nearest nontrivial neighbors* $\texttt{KNTN}()$;

$$x_j \in \texttt{KNTN}(x_i) \Leftrightarrow j = i+1 \text{ or } i \neq j \text{ and } D^l_{x_i,x_j} \leq D^l_{x_i,x_k}, k-\epsilon_w \leq k \leq k+\epsilon_w \quad (1)$$

3) reduce distances in $D^l_{S_i,S_j}$ between points with common (CTN) and adjacent (ATN) temporal relationships:

$$D^0_{x_i,x_j} = \begin{cases} D^l_{x_i,x_j}/(c_{\text{CTN}}c_{\text{ATN}}) & \text{if } x_j \in \texttt{CTN}(x_i) \text{ and } j = i+1 \\ D^l_{x_i,x_j}/c_{\text{CTN}} & \text{if } x_j \in \texttt{CTN}(x_i) \\ D^l_{x_i,x_j}/c_{\text{ATN}} & \text{if } j = i+1 \\ \texttt{penalty}(x_i,x_j) & \text{otherwise} \end{cases} \quad (2)$$

4) transform D^0 into a full all-pairs shortest-path distance matrix, $D = D^g$ (via Dijkstra's algorithm);
5) embed D into d_e-dimensional embedding space through MDS.

where $\texttt{nbhd()}$ are the local neighbors of given segment, c_{CTN} and c_{ATN} are constants for increasing similarity between common and adjacent temporal neighbors. $\texttt{penalty}(x_i,x_j)$ is a function that determines the distance between a pair with no temporal relationship, typically set as $D^l_{x_i,x_j}$.

As in the case with windowed MDS, velocity alone provides only for *proximal disambiguation* of spatially proximal yet structurally different data points in a time-series. By establishing local CTN correspondences, ST-Isomap performs disambiguation. Additionally, these local CTN correspondences are propagated globally through Dijkstra's algorithm to establish datapairs that are *distal correspondences*. Such distal correspondences apply *transitivity* between a datapairs with CTN relationships.

Fig. 2. Robonaut, NASA's space capable humanoid robot.

IV. Experimental Infrastructure

We hypothesize that latent embeddings intrinsically parameterize a robot's sensory-motor space. To explore this hypothesis, we collected collected sets of experimental data (sensory and sensory-motor) from teleoperation trials of the NASA Robonaut.

Robonaut [19] (Fig. 2) was developed by the Dexterous Robotics Laboratory at NASA's Johnson Space Center [1]. Robonaut has two seven degree of freedom (DoF) arms (approximately the size of a human arm), two 12-DoF hands, and a 19-DoF upper extremity. Robonaut's hands have manual dexterity sufficient to perform a wide variety of manipulation tasks. During the course of its operation, Robonaut publishes a 110-dimensional vector time-series at a nominal rate of 50Hz. This time-series data includes information about command torques, proprioceptive sensing (tactile, force, pose), inverse kinematics, and visual exteroceptive sensing.

Although capable of autonomous operation, Robonaut is most typically controlled via teleoperation. The operator is provided the robot's viewpoint with real-time video from Robonaut's eye-centered stereo cameras displayed through an immersive head-mounted display. Sensors in gloves worn by the operator determine Robonaut's finger positions. Electromagnetic motion capture determines the global position and orientation of operator's head and hands. An operator guides the robot using only vision without haptic or force feedback.

A. Data Collection and Robonaut Teleoperation

To explore this hypothesis, we collected collected four sets of experimental data from teleoperation trials of the NASA

Robonaut. The first set consisted of sensory data from various object grasping and manipulation trials. We present results from processing this sensory data as a baseline for evaluating embeddings of sensory-motor data. Sensory-motor data was collected on teleoperation trials for a reach-grasp-lift-move-release task.

As Robonaut operates, either autonomously or via teleoperation, it continually publishes its sensory and motor data at a nominal rate of 50 Hz. The 110 scalar signals were recorded during the experiments and later analyzed as a 110-dimensional vector time-series. The four experiments were:

1) Robonaut reached toward, and grasped a vertically oriented wrench five times at each of nine locations in its workspace.
2) Robonaut held a power drill rigged to be a socket driver, such that it could push the trigger with its index finger. The robot was teleoperated to mate the socket to four lug nuts on a wheel and actuate the trigger to rotate the socket. Four trials were deliberate failures and twenty (five each at each of the four lug nuts) were successful.
3) Robonaut reached for, grasped, and lifted the power drill out of a holster then verified that it was in the correct position to actuate the trigger. This was repeated four times.
4) Robonaut reached for, grasped, lifted, and dropped in another location a vertically oriented chisel. Five of these trials were successful, eight were not.

For experiment 1, five of the teleoperated grasps were analyzed. In experiments 2 and 3, all the trials were analyzed. PCA, windowed MDS, and ST-Isomap were applied to sensory data in each of these experiments. In experiment 4, we analyzed the 5 successful trials as one time series and the 8 unsuccessful trials as another. We analyzed sensory and motor data both collectively and individually, comprising a total of 6 data sets. In all data sets, all the signals were independently mean subtracted then normalized to have unity absolute maxima. The calculations were performed over the extent of the entire time series, not on a per-trial basis.

The results of the 4th experiment indicate that, with normalization, the motor data predominates. There were noticeable differences between the sensory-motor and motor only results but they were not vastly disparate (refer to Fig. 5). While this occurrence does not invalidate our manifold assumptions, it does suggest that more appropriate normalization or sensor fusion methods are warranted.

V. Results

The ST-Isomap embedding of the wrench experiment sensory data is shown in Figure 3 with a comparison to embedding by PCA. From (a) and (b), the structure of the grasps can be surmised from the PCA embedding as two clusters that are transitioned between 5 times. However, this embedding does not clearly identify the procedural similarities in the trials. In contrast, the common procedural structure of the task trials is apparent in the ST-Isomap embedding, (c)-(e). There are five contours that more or less parallel each other. Each represents the sensory response of a single task trial. These appear to trace a clear sensory manifold made up of four (perhaps five) sub-manifolds. Each sub-manifold corresponds to a specific behavior in the task. The black points mark the start of each trial. The green curves (on the right lobe of the "bow-tie" structure in (c) and (d) and on the left in (e)) represent the reach toward the wrench. The red lines (that connect the two lobes of the bow through a short distance) represent the closure of the grasp on the wrench. The cyan contours that make up the bulk of the other lobe indicate the holding of the wrench. The long magenta curves (that connect the lobes through a longer distance) represent the grasp release. The yellow contours (that are intermingled with the green ones) correspond to the retraction of the arm from the wrench.

The structure uncovered from the reach and grasp data appears to characterize the sensor values at different episodes during the task. If that were actually so, then the model would describe the evolution of the sensory response to another instance of the task (that was not used in the original analysis). Given sensory data for this new trial, the training grasps, and the embedded training grasps, interpolation can be used to map the new sensory-data time-series onto the structure found in the embedding space. Shepard's interpolation

The manifolds estimated by ST-Isomap for sensory data from drill mating and holster grasp did not produce as clear a mapping of the sensory data. An appropriately extracted looping structure can be seen in the ST-Isomap embedding. However, this result required a significant amount of parameter tuning. We believe the artifacts in the ST-Isomap embedding are the result of estimating *hard* spatio-temporal correspondences between common temporal neighbors. Such hard decisions with a fixed kernel make selecting a specific setting for neighborhood size a difficult, if not impossible, task. We are currently exploring methods to incorporate soft decisions and adaptive neighboring in spatio-temporal neighborhood determination.

Our inclinations about local neighboring and correspondence are supported when observing the disambiguation results yielded by windowed MDS. Windowed MDS on the drill mating data successfully uncovered the basic looping structure of the mating task. Furthermore, the pressure applied during the mating of the drill and the nut is accurately disambiguated in the embedding. In the holster grasping task, none of the dimension reduction techniques appeared to uncover meaningful structure. We attributed the difficulty in analyzing this dataset to the amount of data being overly sparse for the amount of variation contained in the data.

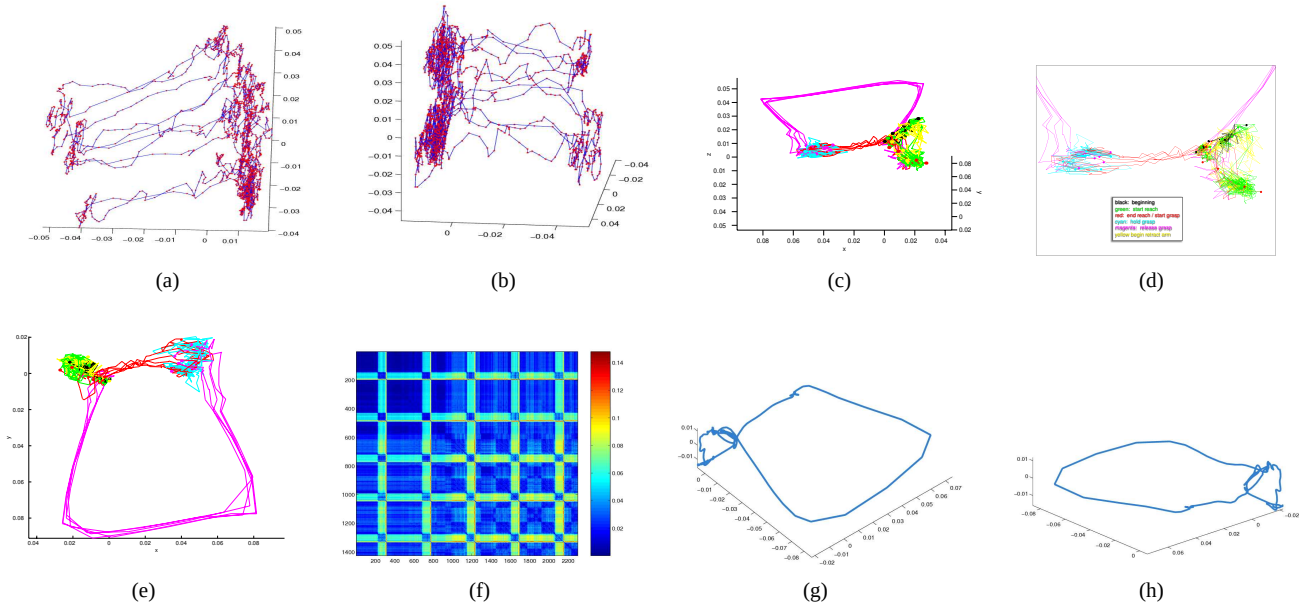

Fig. 3. (a,b) Two views of the PCA embedding for the grasp data from Robonaut teleoperation. (c,d) Two views of the same data embedded by sequentially continuous ST-Isomap. (f) Distance matrix for the ST-Isomap embedding. (g,h) A test grasp mapped via Shepards interpolation onto the grasp structure in the ST-Isomap embedding.

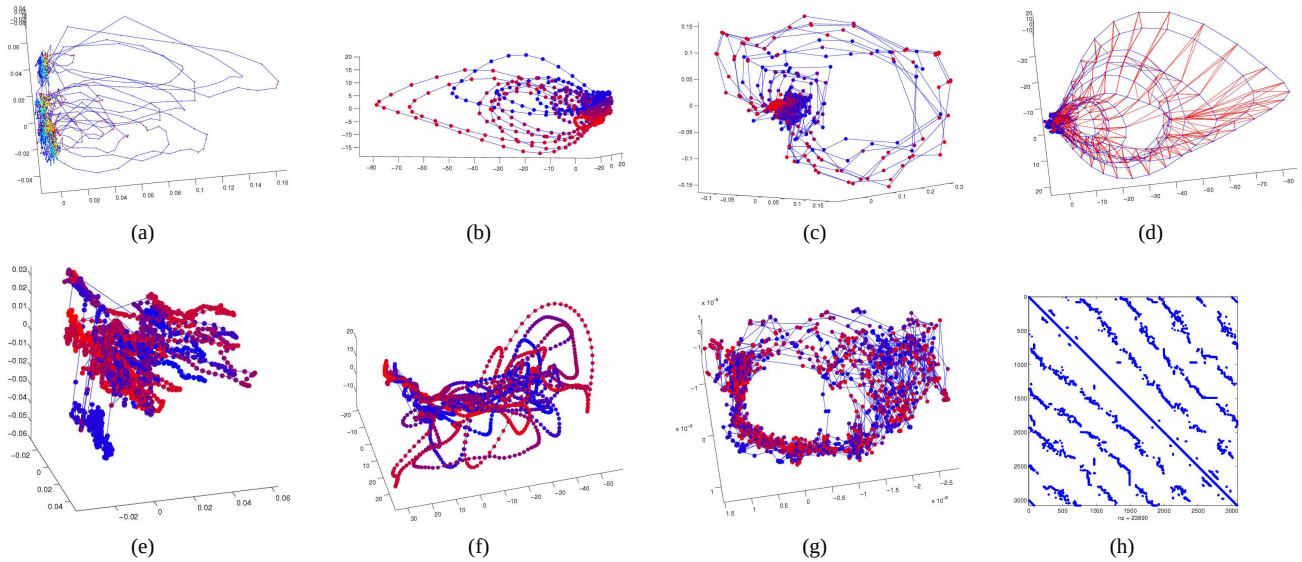

Fig. 4. Embedding of the (top) drill-to-nut mating and (bottom) take-drill-out-of-holster task sensor data sets with (a,e) PCA, (b,f) windowed MDS, and (c,g) ST-Isomap. (d) Spatio-temporal correspondences computed by ST-Isomap for the mating task. (h) Spatio-temporal distance matrix for the holster task.

VI. Conclusion

Our analysis of the sensory data from Robonaut yielded excellent structural results in one of the three cases analyzed, with mixed results in other sensory analysis experiments. Our initial explorations into sensory-motor manifold learning demonstrated some promise with the uncovering of visually observable clusters. While promising, several questions for manifold learning with regard to neighborhood determination, spatio-temporal correspondence, and sensor modality normalization remain further questions to fully answer our manifold hypothesis.

Acknowledgements

This work was funded in part by NSF Award IIS-0534858 and NASA grants NAG9-1446, NAG9-1515, and

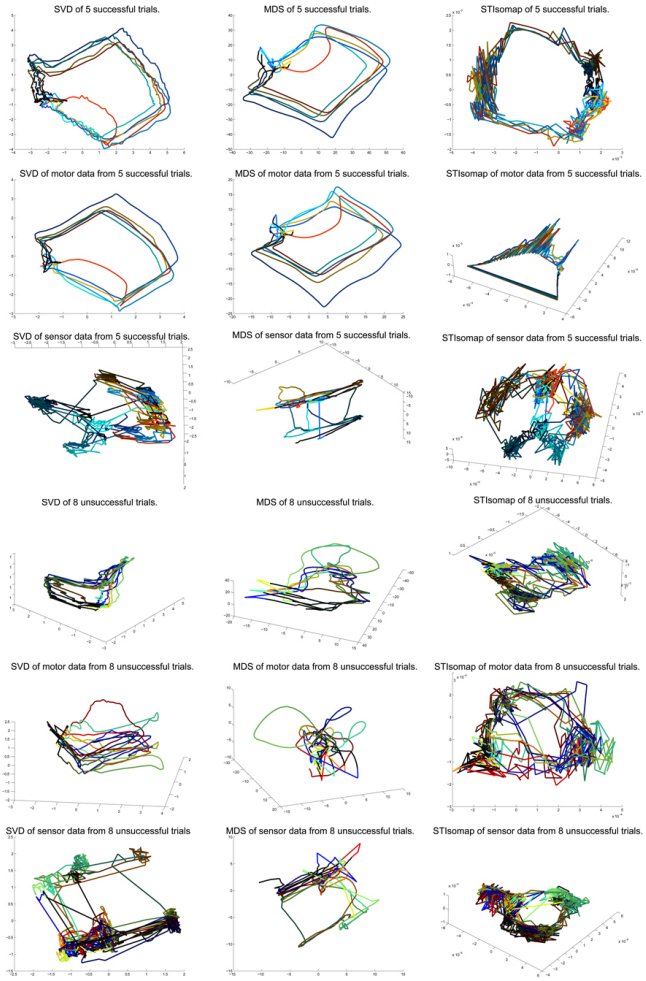

Fig. 5. Embeddings of grasp-drop sensory-motor data PCA (column-nwise left), windowed MDS (middle), ST-Isomap (right). Each row respectively contains sensory-motor, sensory, and motor data from successful (top three rows) and unsuccessful trials (bottom three rows).

NNJ04HI19G. We thank the Robonaut Team at Johnson Space Center for their assistance during data collection.

REFERENCES

[1] R. O. Ambrose, H. Aldridge, R. S. Askew, R. R. Burridge, W. Bluethmann, M. Diftler, C. Lovchik, D. Magruder, and F. Rehnmark. Robonaut: Nasa's space humanoid. *IEEE Intelligent Systems*, 15(4):57–63, July 2000.

[2] M. Asada, M. Ogino, S. Matsuyama, and J. Ooga. Imitation learning based on visuo-somatic mapping. In *Proceedings of 9th International Symposium on Experimental Robotics*, volume CD-ROM, Singapore, June 2004. Springer-Verlag.

[3] C. G. Atkeson, J. G. Hale, F. Pollick, M. J. Riley, S. Kotosaka, S. Schaal, T. Shibata, G. Tevatia, A. Ude, S. Vijayakumar, and M. Kawato. Using humanoid robots to study human behavior. *IEEE Intelligent Systems*, 15(4):46–56, July 2000.

[4] C. G. Atkeson and S. Schaal. Robot learning from demonstration. In *ICML '97: Proceedings of the Fourteenth International Conference on Machine Learning*, pages 12–20, San Francisco, CA, USA, 1997. Morgan Kaufmann Publishers Inc.

[5] M. E. Cambron and R. A. Peters II. Determination of sensory motor coordination parameters for a robot via teleoperation. In *Proceedings of the 2001 IEEE International Conference on Systems, Man, and Cybernetics*, volume 5, pages 3252–3257, Oct. 2001.

[6] C. L. Campbell, R. A. Peters II, R. E. Bodenheimer, W. J. Bluethmann, E. Huber, and R. O. Ambrose. Superpositioning of behaviors learned through teleoperation. *IEEE Transactions on Robotics*, to appear, 2006.

[7] J. A. Coelho Jr., J. H. Piater, and R. A. Grupen. Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot. In *Proceedings of the First IEEE/RAS International conference on Humanoid Robots (Humanoids 2000)*, Sept. 2000.

[8] P. R. Cohen and N. Adams. An algorithm for segmenting categorical time series into meaningful episodes. In *Proceedings of the Fourth Symposium on Intelligent Data Analysis*, volume 2189, pages 198–207, 2001.

[9] T. Cox and M. Cox. *Multidimensional Scaling*. Chapman and Hall, London, 1994.

[10] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification (2nd Edition)*. Wiley-Interscience, 2000.

[11] O. C. Jenkins and M. J. Matarić. Performance-derived behavior vocabularies: data-driven acqusition of skills from motion. *International Journal of Humanoid Robotics*, 1(2):237–288, June 2004.

[12] O. C. Jenkins and M. J. Matarić. A spatio-temporal extension to isomap nonlinear dimension reduction. In *The International Conference on Machine Learning (ICML 2004)*, pages 441–448, Banff, Alberta, Canada, July 2004.

[13] M. Lungarella, T. Pegors, D. Bulwinkle, and O. Sporns. Methods for quantifying the information structure of sensory and motor data. *Neuroinformatics*, 3(3):243–262, Fall 2005.

[14] K. F. MacDorman, R. Chalodhorn, and M. Asada. Periodic nonlinear principal component neural networks for humanoid motion segmentation, generalization, and generation. In *Proceedings of the Seventeenth International Conference on Pattern Recognition*, pages 537–540, Cambridge, UK, Aug. 2004. International Association for Pattern Recognition.

[15] M. J. Matarić. Getting humanoids to move and imitate. *IEEE Intelligent Systems*, 15(4):18–24, July 2000.

[16] R. A. Peters II, C. L. Campbell, W. J. Bluethmann, and E. Huber. Robonaut task learning through teleoperation. In *Proceedings of the 2003 IEEE International Conference on Robots and Automation*, Taipei, Taiwan, Oct. 2003.

[17] R. Pfeifer and C. Scheier. *Understanding Intelligence*. The MIT Press, Cambridge, MA, 1999.

[18] N. S. Pollard, J. K. Hodgins, M. J. Riley, and C. G. Atkeson. Adapting human motion for the control of a humanoid robot. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, pages 1390 – 1397, May 2002.

[19] Robonaut Development Team. Robonaut avionics. Technical report, NASA Johnson Space Center, http://robonaut.jsc.nasa.gov/Avionics.htm, July 2004.

[20] S. Roweis and L. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, Dec. 2000.

[21] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[22] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290:2319–2323, 22 December 2000.

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

[24] A. Ude, C. G. Atkeson, and M. J. Riley. Planning of joint trajectories for humanoid robots using b-spline wavelets. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2223–2228, Apr. 2000.

[25] C. K. I. Williams. On a connection between kernel pca and metric multidimensional scaling. *Machine Learning*, 46(1-3):11–19, 2002.