Uncovering Success in Manipulation

Odest Chadwicke Jenkins
Department of Computer Science
Brown University
Providence, RI 02912-1910
Email: cjenkins@cs.brown.edu

Richard Alan Peters II
Center for Intelligent Systems
Vanderbilt University School of Engineering
Nashville, TN 37235
Email: Alan.Peters@Vanderbilt.Edu

Robert E. Bodenheimer
Center for Intelligent Systems
Vanderbilt University School of Engineering
Nashville, TN 37235
Email: bobbyb@vuse.vanderbilt.edu

***Abstract*— Experiments were performed with the NASA Robonaut to determine if manifold learning could discern successful and unsuccessful teleoperation trials in an unsupervised manner. Repeated teleoperation of drill-mating and chisel-pickup tasks were performed by a skilled teleoperator. Spatio-temporal Isomap (STI) was used to embed data from the robot's sensory-motor state-space (SMSS) to uncover underlying structure and separability between successful and unsuccessful trials. We present results from embedding SMSS data from repeated teleoperation performances, visualized in 3 dimensions, where success and unsuccessful trials are discerned. Our results are further evaluated by out-of-sample projection and comparison with Support Vector Machines for classifying the success of new teleoperation trials.**

I. Introduction

As robots move from industrial floors to domestic settings, the need for natural methods for human-robot collaboration becomes crucial for successful deployment. In particular, robot programming by demonstration, rather than computer programming, is increasingly important for the transfer of skills from human users to robots. However, humans have an innate sense for goal attainment and sensorimotor control that robots lack. Our research objectives involve endowing robots with the ability to acquire sensorimotor skills from human demonstration. Such sensorimotor skills must be able to function in various environments, perform in the face of uncertainty due to partial observations, and have a sense of its interactions with the environment. We phrase these low-level sensorimotor issues as a learning problem. That is, given a specific skill of interest, can the structure underlying successful performance of skills and task be uncovered as a manifold in the robot's sensorimotor space? If this question can be answered, suitable detection mechanisms can be developed for detecting and recovering from unsuccessful robot execution. Additionally, such endeavors will provide an avenue for learning robust manifold attractors, along the lines of [1], [2], from demonstration.

When a robot is programmed through demonstration or controlled through teleoperation, its resultant sensory-motor data stream can form discernable patterns in the vector space that contains them, the sensory-motor state space (SMSS). The patterns reflect both measurable effects on the environment of the robot's actions and its motor reactions to sensory input. Thus, the patterns emerge from a closed-loop interaction between robot and environment. This phenomenon was demonstrated with a simple mobile manipulation robot by Pfeifer in 1999 [3]. The SMSS has dimension equal to the number of scalar signals that can be recorded while the robot operates. But, the effective dimension of the pattern may be much smaller, depending on the number of independent variables that dominate during the interaction. In cases of repetitive, constrained motion by the robot (for example repeatedly reaching toward and grasping an object) the dominant variables tend to trace closed manifolds in the SMSS. Closure makes sense because during exact repetitions of a task the trajectory through SMSS would repeat itself. If the task is repeated with some variations, say under different initial conditions in robot or environment, the trajectory does not repeat itself exactly. Instead, a family of trajectories lies on a manifold in the SMSS, displaced from one another along directions that correspond to the variations. By having the robot perform the same task under different initial conditions, limits on the manifold might be discerned.

This paper reports on two sets of results designed to elicit a bifurcated manifold and to determine if it could be used to to classify further repetitions of the task. The experiments, performed with the NASA Robonaut [4], was to: 1) reach for, grasp, pick up, move, and release an object, then return to the starting position and 2) mate a drill socket to a nut on a wheel. The teleoperator caused the robot to succeed during some tasks and fail during others. To determine if the SMSS vectors in the recorded data could be classified with a probability greater than chance, a Support Vector Machine (SVM) Analysis was used. Since the task had two possible outcomes over quasi-periodic repetitions the dominant patterns in the SMSS should be low dimensional – at least 2D, perhaps 3D. To elucidate any such manifolds, manifold learning in the form of spatio-temporal Isomap [5] was applied to the sensorimotor time-series.

II. Previous Work

In [6] 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 [7]. The formation of low dimensional manifolds in the Robonaut SMSS as a consequence of task repetition was reported in [8]. In addition to Pfeifer [3], many others have studied the extraction of SMC parameters.

Fig. 1. Robonaut, NASA's space capable humanoid robot.

Jenkins and Matarić developed Spatio-temporal Isomap (STI) for the creation of new motions through the interpolation of learned trajectories [9]. STI is an extension of Isomap [10], one of a number of dimensionality reduction techniques including Principal Component Analysis [11], and the related technique of Multi-dimensional Scaling [12].

Support vector machines are described in several textbooks including [13]. Pelossof, et al., [14] studied the learning of stable grasps by SVMs.

III. Robonaut

Robonaut [4] is NASA's space-capable, humanoid robot (Fig. 1), developed by its Dexterous Robotics Laboratory. Each seven degree of freedom (DoF) Robonaut arm is approximately the size of a human arm. Each of those mates with a 12-DoF hand to produce a 19-DoF upper extremity.

Robonaut's sensors include two hand/wrist modules, containing 98 sensors for feedback and control. Each DoF has a motor position sensor, a joint force sensor, and a joint absolute position sensor. The two arm modules contain 90 sensors. Each actuator contains a motor incremental position sensor, redundant joint torque sensors, redundant joint absolute position sensors, and four temperature sensors distributed throughout the joint. Each arm employs relative optical encoders in five of its joints. The encoders reside on the motor side of the gear train and have resolutions ranging between 200 and 1000 counts per degree of arm motion.

The data signals that were recorded from Robonaut during teleoperation are listed in Table I. The ones that were actually used for this experiment are in italic type. The resulting 105-dimensional vector time-series was recorded at a nominal rate of 8Hz.

Although Robonaut is physically capable of autonomous operation it is most often controlled directly via teleoperation. Significantly, haptic sensations and joint forces cannot be reflected from the robot back to the teleoperator, who guides the robot based on vision alone. Given its sensor suite, the robot is capable of "feeling" for itself the effects of its actions.

TABLE I

Signals Recorded from Robonaut.

Signal	Dimension
End-effector 4×4 position	16
Arm orbit angle	1
Arm joint positions	7
Finger joint positions	12
6-axis force on wrist	6
6-axis force on shoulder	6
Arm joint torques	7
Force on fingers	5
Finger joint torques	12
Hand tactile sensors	33

To enable the robot to act and react on its own sensory motor coordination is one of the motivations behind the research reported herein.

IV. Manifolds in SMSS and Outcome Classification

In [8] it was reported that in Robonaut's SMSS closed manifolds can be formed by task repetition. If the robot always starts the task in a similar SM state, the initial part of the manifold should be small, highly localized. If the task diverges into multiple variations or outcomes, one would expect the manifold to ramify accordingly over the course of the repetitions. In particular, if the task has the binary outcome set $\{\text{success}, \text{failure}\}$, and if the trials are performed to include examples of each, one would expect the manifold to bifurcate. If this were shown to be true, later repetitions of the task should map to the manifold in such a way that success or failure could be detected. That is, by learning the manifold under teleoperation, the robot could then assess the outcome of later autonomous execution by observing the branch of the manifold upon which its SM state projects.

A. Support Vector Machines

The SVM algorithm operates by mapping a given training set into a high-dimensional *feature space* and finding a hyperplane that separates the data into classes. To construct an optimal hyperplane, the SVM minimizes a particular error function, and in this work, we use the C-SVM classification [15]. Given a training set of attribute-label pairs (x_i, y_i), where $i = 1 \ldots l$, training vectors $x_i \in \mathbb{R}^n$ and $y_i \in \{+1, -1\}^l$, C-SVM minimizes the following error function:

$$\min_{\mathbf{w},b,\xi} \frac{1}{2}\mathbf{w}^T\mathbf{w} + C\sum_{i=1}^{l}\xi_i$$

subject to $y_i(\mathbf{w}^T\phi(x_i) + b) \geq 1 - \xi_i$. The training vectors x_i are mapped to a higher dimension by the kernel function ϕ. Given a sufficiently high dimension and an appropriate, nonlinear kernel ϕ, any data set can be mapped by ϕ into the high dimensional space such that a hyperplane separating the data into its appropriate categories exists. C is the penalty parameter of the error function, which controls the trade-off between allowing training errors and forcing rigid margins, $\mathbf{w}$

is a vector of coefficients, b is a constant, and ξ_i are parameters for handling non-separable input data. We chose to use a radial basis function (RBF) kernel having the form $\phi = e^{-\gamma\|x_i - x_j\|^2}$, where $\gamma > 0$.

B. Manifold Learning through Dimensionality Reduction

We assume sensorimotor observables – the time series – are intrinsically parameterized by a lower dimensional embedding. The embedding provides a mapping $\mathbf{x} = \phi(\mathbf{y})$ between intrinsic parameters and observations, realizing intrinsic coordinates $\mathbf{y}_k : \{1, 2, 3, \ldots N\} \rightarrow \mathbb{R}^n$ for the input data where $n < N$. Such a latent parameterization could be uncovered by applying dimensionality reduction techniques such as Principal Components Analysis [11]. PCA involves an eigendecomposition on a linear covariance matrix to find an orthogonal subspace of principal components that compactly approximate the input data. Singular value decomposition, which we have used, provides an equivalent projection.

Multidimensional scaling [12] is another approach where pairwise distances, rather than linear covariance, are preserved. Given the distance between all input data pairs $D_{\mathbf{s}_j, \mathbf{s}_k}$, MDS produces embedding coordinates that minimizes the error $E = |D_s - D_y|_{L^2}$, where D_s and D_y are respectively the pairwise input and embedding space distance matrices. Essentially, MDS produces embedding coordinates that preserves the distance metric as much as possible. Isomap [10] uses a geodesic (Dijkstra shortest-path) distance metric with MDS. This technique can be summarized by a three step process: 1) finding the nearest neighorbors of each point forming a sparse pairwise distance matrix, 2) filling this distance matrix using Dijkstra shortest-path computation from each datapoint, 3) embedding of the full distance matrix into d dimensional coordinates through MDS. The resulting embeddings avoid "short-circuiting" problems associated with Euclidean distance between non-proximal data pairs.

These techniques are not ideally suited for time-series analysis data because they assume the input are i.i.d. – independent samples from the same manifold parameterization. Time-series data are not independent, but rather sequentially ordered samples from an underlying spatio-temporal process. To add a time-dependency to MDS, we use a "windowed MDS" procedure, where each input data object is a temporally extended window of observations. Such windows extend over a fixed horizon of time. Adding time as another dimension This serves to disambiguate spatially proximal data pairs that are different phases of a temporal process. But it does not detect the temporal coincidence of spatially distant data pairs that are in phase with respect to the temporal process.

Spatio-temporal Isomap (STI) [16] combines the consideration of temporality with the geodesic similarity propagation of Isomap. This method follows the format of Isomap except that the nearest neighbors are spatio-temporal. Spatio-temporal neighbors to a given window are the closest windows, given by L^2 distance, that are not trivially related to better matching windows. Two data windows are consider trivially related if they occur close in time (within some ϵ threshold). STI provides both the ability to disambiguate data for classification and correspond for registration and clustering.

Fig. 2. Robonaut drill mating sequence.

V. The Experiments

We performed two experiments to evaluate the suitablity for STI to uncover structure in teleoperated manipulation trials where success is mixed. For the drill-mating task, an embedding was formed using STI of equal successful and unsuccessful trials. Test trials were projected into this embedding to visualize whether they followed uncovered structures for success or failure. In the chisel-pickup task, data from mixed success trials were formed into an STI embedding and compared against training an SVM on a subset of the trials. This experiment compares the supervised SVM, which requires training labels, against the unsupervised STI, which can embed training and test data together because labels are not required.

A. The Drill-mating Task

In drill mating task (Fig. 2), Robonaut was teleoperated to mate the socket of a drill it was holding with a nut on a statically positioned wheel. Eight trials of the task were performed. 4 of the 8 trials the task were completed successfully. The unsuccessful trials resulted in the some collision of the socket and nut without closure. SMSS data from six of the trials (3 successful and 3 failure) were embedded using STI. As shown in panels (a) and (b) Figure 3, STI embedding was able to discern the two classes of success in the teleoperation sensorimotor data. In addition, STI was able to uncover temporal regularity in the successful trials in the form of a "looping" structure. We interpret this loop as the registration of a spatio-temporal signature common to all of the successful trials.

The classification capability of the STI manifold embedding was evaluated through out-of-sample projection of test trials. The two remaining trials from teleoperation were projected into the embedding using Shepard's interpolation. Shepard's interpolation is a technique that reconstructs a data point based on the average of other points in the data set weighted by distance. For STI projections, the distance weight is computed using L^2 between window horizons. Panels (d) and (e) show the projection the test successful and failure trials, respectively. From manual observation, it is clear that these test trials are appropriately discerned based on the actual success of the trial. The successful test trial does conform to the same looping structure as in the other successful trials. However, this appears to somewhat of a lesser degree, which we attribute the light density of training trials and simplistic nature of Shepard's interpolation.

B. The Chisel-pickup Task

Robonaut was teleoperated through a task that involved reaching, grasping, and moving an object. Thirteen trials of the task were performed. In 5 of the trials the task was completed successfully; in 8, it was not (Fig. 4). The object was an upright chisel on a stand. A simple distal closure grasp could be used. A trial was a success if the robot formed a stable grasp on the object, lifted it, moved it to another position where it released it. A trial was a failure if the robot knocked the object over without forming a grasp or if the robot let slip the object upon lifting it.

!!! The 3D motion trajectories of the end-effector are shown in Fig. 6 (a). The successful trials are in blue and the unsuccessful in red. The starting point of all trials was near $(x, y) = (38, 42)$. The chisel was randomly placed near $(x, y) = (58, 42)$. There was more variation in the object placement than the starting position. In every successful trial the robot moved the object to a position near $(x, y) = (38, 72)$ where it released it and returned to the starting area in a predominately y-axis direction. In all the unsuccessful trials the robot moved its end-effector beyond the object position and returned in a predominately y-axis direction.

It is clear from the figure that if end-effector position or arm joint angles were used to train the manifold, they would dominate its structure and make classification by any method trivial. Less obvious from the diagram, but quite evident in the motor signals, the wrist position differed significantly between the two classes on a segment of the trajectories. We excluded that information from the analysis and used only end-effector related signals. Thus classification of the outcome was dependent solely on sensory information from the wrist to the finger tips and on motor information from the fingers alone (Table I). All 110 sensory and motor channels were sampled at 8Hz. The exclusion of arm-related signals left a 68-dimensional vector time series. The number of vectors in a trial varied from 100 to 215 with a median length of 161 vectors.

Support vector machine analysis was used to estimate the probability that any single vector from any single trial could be identified correctly as coming from a successful trial. An SVM classifier was built from a subset of the trials then used to classify the vectors in the remaining trials. The analysis was performed using motor data alone, sensory data alone, and sensory and motor data together. A radial basis function (RBF) kernel was used for all the tests. The associated parameters, margin $C = 32$ and RBF exponent $\gamma = 8$ were determined through a grid search over $C = 2^{-5}, 2^3, \ldots, 2^{15}$ and $\gamma = 2^{-15}, 2^{-13}, \ldots, 2^5$ using $4\times$ cross-validation as suggested in [17].

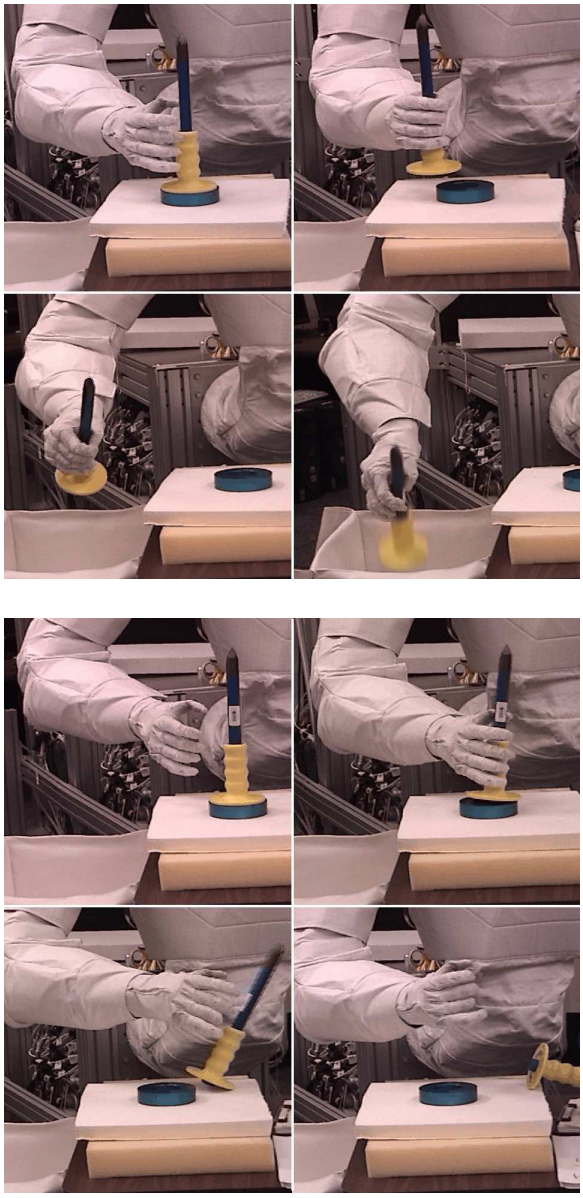

Fig. 4. Successful (top) and unsuccessful (bottom) grasp sequences.

The time-series of the 13 trials were analyzed using Singular Value Decomposition, Multidimensional Scaling, and Spatio-Temporal Isomap. The three most significant dimensions of contours were plotted such that the first two principal direc-

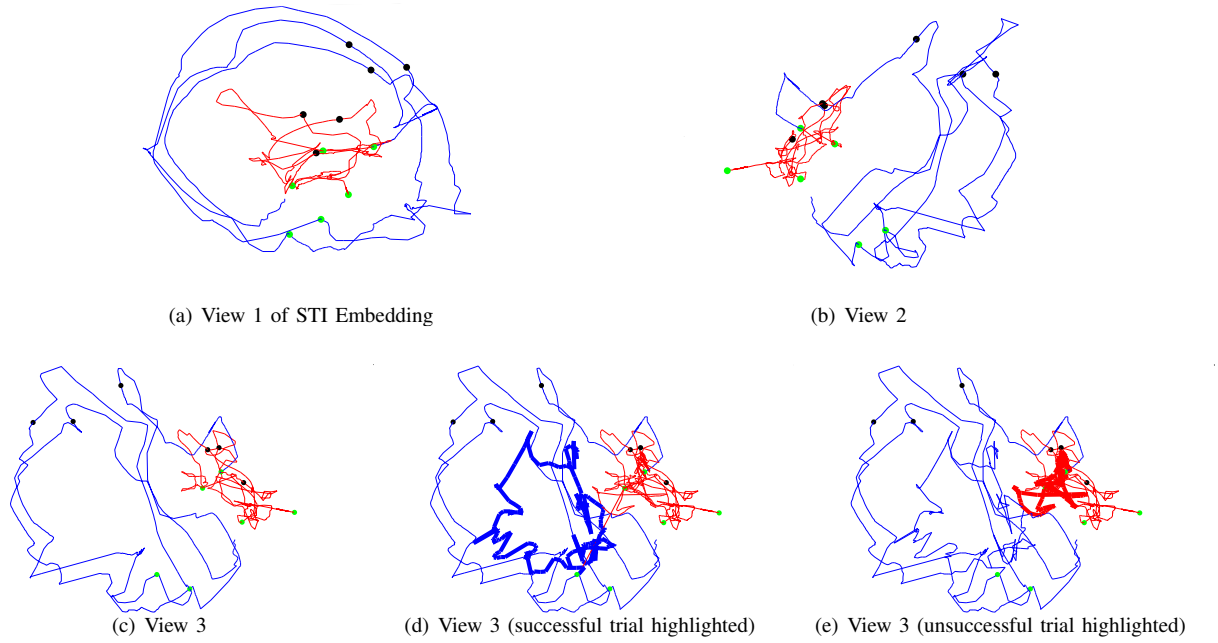

(a) View 1 of STI Embedding

(b) View 2

(c) View 3

(d) View 3 (successful trial highlighted)

(e) View 3 (unsuccessful trial highlighted)

Fig. 3. (a)-(c) Embedding of 3 successful (in blue) and 3 unsuccessful (in red) drill mating trials viewed from three different viewpoints. Out-of-sample projection of new (d) successful and (e) unsuccessful trials are shown in bold.

TABLE II
PERCENT VECTORS CORRECTLY CLASSIFIED BY SVM

Sensory Data	Motor	Sensory	Sensory-Motor
Unprocessed	69	64	63
Processed	69	67	70

tions defined the xy-plane. If the SMSS manifold were to exhibit a significant bifurcation, it should be evident in that plane.

1) Results: Table II displays the aggregate results of the exhaustive SVM tests for classifiers built from 2 successful and 4 unsuccessful trials applied to the unprocessed and the processed data, further broken down into motor signals alone, sensory signals alone, and sensory motor together. When using the data directly from the robot (normalized but otherwise unprocessed), the motor data alone yields a better classifier than either the sensory data or the sensory-motor together. Applying the nonlinear noise filter to the haptic signals resulted in the sensory-motor classifier being the best. But was only slightly better. Moreover a correct classification of only 70% is not particularly good. These probabilities are the averages over the 700 models built from 2 successful and 4 unsuccessful trials. The best model of that type correctly classified the processed sensory-motor data with 87% accuracy and the worst with 44%.

To determine which vectors were being misclassified, we selected a single classifier, the one trained on trials 1, 2, 4, 5, 10 and 13. Trials 2 and 4 were successful. The other four were not. That classifier had 73% accuracy when applied to the vectors in the 7 trials not used for model, *viz.* 3, 6, 7, 8, 9, 11, and 12, with trials 3 and 11 the successful ones. Fig. 6 (a) shows the misclassified vectors plotted as '$*$'s on the end-effector trajectory. Many of those points were along trajectory 3 for reasons unknown.

We applied the 3 dimensionality reduction procedures to all 13 trials in one continuous time series. The top row of Fig. 5 shows results for the full 68-D unprocessed time series. The bottom row shows them for the 19-D processed time series. The first panel, (a), depicts the windowed distance matrix from which the MDS and STI embeddings were computed. Panel (b) shows the SVD, (c) the windowed MDS, and (d) the STI embedding of the time series. All 6 of these clearly bifurcate along task outcome in the principal plane. The manifold traced by the STI embedding of the processed data forms has the best separation and the most symmetric structure.

Fig. 6 (b) and (c) compare the SVM classification to classification by STI. Panel (b) is a plot of the SVM misclassified vectors on the STI embedding. Most of the misclassifications occur in the reach phases of the tasks as would be expected, given causality. But the misclassification of trial 3's post grasp trajectory is also visible.

The thin contours in Fig. 6 (c) comprise the STI embedding of trials 1, 2, 4, 5, 10 and 13 – those used to train the SVM. The thick lines show the projection of trials 3, 6, 7, 8, 9, 11, and 12, onto the manifold traced by the former. This shows that the manifold embedding created by STI using 6 trials of the task not only traces an outcome-dependent manifold, but also classifies the new data much more accurately than does the SVM classifier.

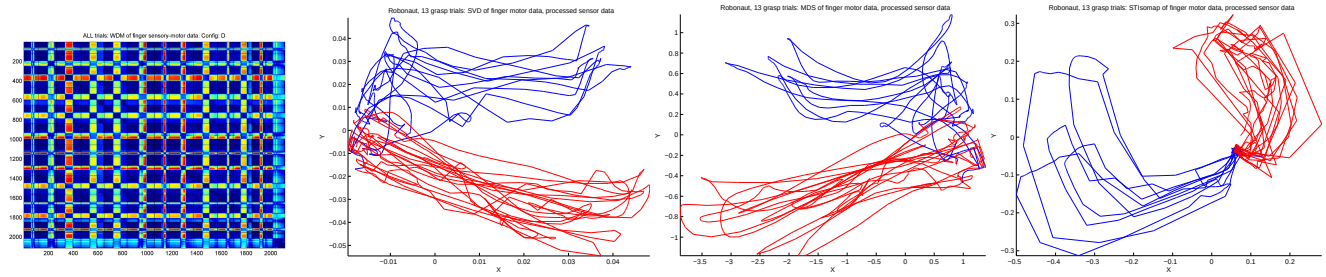

(a) (b) (c) (d)

Fig. 5. In all the trajectory plots, blue corresponds to a successful, and red to an unsuccessful, trial. (a) Windowed distance matrix of the sensory-motor time series. (b) Trajectory embedded by SVD-PCA, (c) by MDS, and (d) by STIsomap.

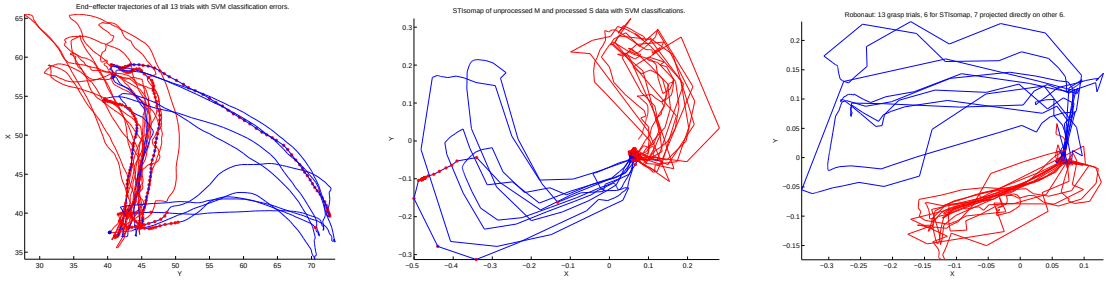

(a) (b) (c)

Fig. 6. In (a)-(b) a red * is the location of a point on a successful trajectory that was misclassified by the SVM and blue star is an incorrectly classified point on an unsuccessful trajectory. (b)-(c) were generated from the same training and test set of vectors. (a) vectors misclassified by SVM plotted on the end-effector (real-world) trajectory. (b) Vectors misclassified by SVM plotted on STI Trajectory (Unprocessed motor, processed sensor). (c) Projection of 3 successful and 4 unsuccessful trajectories (thick lines) onto the STI embedding of 2 successful and 4 unsuccessful trajectories (thin lines).

REFERENCES

[1] R. Platt, Jr., A. H. Fagg, and R. A. Grupen, "Manipulation gaits: sequences of grasp control tasks," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation (ICRA 2004)*, vol. 1, Apr. 2004, pp. 801 – 806.

[2] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential composition of dynamically dexterous robot behaviors," *International Journal of Robotics Research*, vol. 18, no. 6, pp. 534–555, 1999.

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

[4] 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*, vol. 15, no. 4, pp. 57–63, July 2000.

[5] O. C. Jenkins and M. J. Mataríc, "A spatio-temporal extension to isomap nonlinear dimension reduction," in *International Conference on Machine Learning (ICML 2004)*, Banff, Alberta, Canada, July 2004, pp. 441–448.

[6] 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 (ICRA 2003)*, Taipei, Taiwan, Oct. 2003.

[7] 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*, vol. 22, no. 1, pp. 1–13, Feb. 2006.

[8] R. A. Peters II and O. C. Jenkins, "Robonaut: manifold structures in sensory motor state space," in *Proceedings of the IEEE/RAS International conference on Humanoid Robots (Humanoids 2005)*, Tsukuba, Japan, Dec. 2005.

[9] O. C. Jenkins, "Data-driven derivation of skills for autonomous humanoid agents," Ph.D. dissertation, University of Southern California, Robotics Research Laboratory, Center for Robotics and Embedded Systems, Computer Science Department, University of Southern California, 941 W. 37th Place, Los Angeles, CA 90089 USA, 2003.

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

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

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

[13] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines and other kernel-base learning methods*. Cambridge University Press, 2000.

[14] R. Pelossof, A. Miller, P. Allen, and T. Jebara, "An svm learning approach to robotic grasping," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation (ICRA 2004)*, Apr. 2004, pp. 3212–3218.

[15] V. N. Vapnik, *Statistical Learning Theory*. New York: Wiley, 1998.

[16] 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)*, Banff, Alberta, Canada, July 2004, pp. 441–448. [Online]. Available: http://www.cs.brown.edu/~cjenkins/papers/cjenkins_stisomap.pdf

[17] C.-W. Hsu, C.-C. Chang, and C.-J. Lin, "A practical guide to support vector classification," Department of Computer Science, National Taiwan University, Tech. Rep., 2003. [Online]. Available: http://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf