Learning to Detect Spatial Regions with Constraints

Veselin Georgiev, Todd Wegter, Ramy Sweidan, Vladimir Sukhoy, and Alexander Stoytchev
Developmental Robotics Laboratory
Iowa State University
vvgeorgiev2s@semo.edu, twegter@iastate.edu, rms2@rice.edu, {sukhoy, alexs}@iastate.edu

Abstract— **This paper proposes a method for constraint detection using only proprioceptive data. This method for constraint detection was tested on a pre-existing data set and was found to be able to detect insertions. Two additional experiments were performed. The first showed that the constraint detection algorithm can be used to associate an action with an event. The second experiment implemented a learning technique to show that learning from constraint detection can improve the efficiency of performing a task. Our method can be applied to both environmental constraints (i.e. limited space to move) and object constraints (i.e. levers, wheels, etc.).**

I. Introduction

Traditionally, robots trying to complete tasks in constrained space are provided with a pre-defined representation of that space. However, variation between spaces makes it nearly impossible to program a robot to interact with every constraint it could encounter. We propose that this is not the best approach because it limits the robot to completing tasks only in pre-defined spaces. Robots that need to complete tasks in constrained space should be able to explore the constrained space on their own and learn from the exploration. For example, a housekeeping robot would need to unlock and open doors. Trying to program a housekeeping robot to use every type of door handle and lock would be nearly impossible.

We propose a new method which allows a robot to learn how to detect and interact with the constraints in its environment. This allows the robot to explore any environment independently of any other senses such as vision. A learning algorithm is then employed to enable the robot to concentrate on unconstrained movements which augments the robot's ability to complete a task in a constrained space.

II. Related Work

A. Robotics

The current research is highly motivated by the prior work of Koonce et al. [8] who developed a method for a robot to insert a peg into a hole. This method was simple and involved few steps. First, the robot pressed the block up against the aperture in which various shaped holes were cut. Then, the robot rotated the block without observing it until the the block fit in the hole. This experiment successfully used random exploratory behaviors to complete a peg in a hole task without using any kind of sensory feedback. However, because of the random nature of the task and the lack of sensory feedback, the robot could not truly be said to know what exactly it was doing and could not make corrections to its actions should its senses call for them.

Upon examining the proprioceptive data, it was concluded that the propioceptive data for successful insertion trials appeared differed from that of unsuccessful ones. This could possibly be used to give the robot some amount of awareness that its arm is operating within a constrained space and has therefore completed the task. While our task maintains the same basic principles of using exploratory behaviors to perform an insertion task, we attempted to program the robot to learn if its hand is constrained by analyzing joint torques.

While the current research does not explicitly concern insertion or assembly tasks, these types of activities do involve robotic manipulators planning motions around a series of constraints. The completion of these tasks is achieved by a variety of means. One way is by using visual models to guide insertions. Morrow et al. [13] developed visual primitives and combined them with force based primitives to create a sensorimotor algorithm that allowed the robot to complete connector insertion tasks. Meeussen et al. [12] programmed their robot to use laser scanners to find places in a wall that afforded the insertion of a three pronged plug. The robot could then plug itself in, allowing it to act self-sufficiently for long periods of time. Mayton et al. [11] performed a similar feat, in that their robot was able to plug itself in. However, instead of using visual feedback to do so it used the natural electromagnetic fields created by electrical outlets to guide it and to help it insert its plug into the outlet.

Bruyninckx et al. [1] demonstrated that it was possible to use either of two different models in order to complete insertion tasks. The first was a kinematic model that created a model of the goal location (the hole) and calculated three points of constraints. The robot's manipulator would not leave these constrained areas. It would then calculate the Jacobians of the manipulators and the peg and use them to guide its manipulator to the goal. A geometric model was also used, in which the robot took into account the angles and dimensions of the peg and the hole and would arrange the peg so that it would arrive at the hole already angled for entry.

Suarez et al. [18] aimed to find a way to perform insertion tasks without the use of geometric models, as those are not always available in an accurate form. To do this, they modeled the many uncertainties of the robot's environment and used force feedback information to guide the robot's fine motions through several task states and eventually into the goal state. Paetsch and his colleagues [15]circumvented the problem of uncertainties by programming into their robot a number of strategies modeled after real humans' strategies for inserting a peg into a hole. As the robot attempted to place the peg into the hole, it used force feedback sensors to decide which strategies to use. Paetsch's results showed that his strategies, when used in tandem with each other, increased the success rate of the task.

B. Developmental Psychology

As the current research aims to develop the robot's capabilities in an incremental manner, it is important to model it after the growth of an infant's capabilities. Infants begin life with the ability to perform very basic exploratory behaviors such as grasping and pulling objects closer to them. While their movements seem random, they must have some amount of intention behind them as infants express expectations of the results of their movements [4]. They use most of their senses, including touch, taste, sight, and proprioception to discover the affordances of objects. Affordances are what an actor perceives an object's use to be as well as how the object conveys this information. Affordances are subjective and context based; a stick may afford being used as a poking device or it may be used to pull objects close when they are far away[19]. As infants grow, they are able to discern affordances more easily by using perception [9][2]. They can perceive affordances, observe the effects their actions have on their environments, and adjust their ideas of the present affordances [4].

Studying the progression of infants' abilities to manipulate objects in constrained space is directly beneficial to our research. Infants have shown the ability to consistently insert objects into holes at approximately 22 months of age [17]. Infant chimpanzees, which mirror the developmental cycle of humans, exhibit the insertion behaviors even earlier at 10 months [6]. This ability does not simply pop into existence; it follows a series of changes. At 15 months, human infants are able to perform insertion tasks, albeit at rates of chance [14]. This shows that as motor skills develop and cognitive faculties grow, insertion tasks become easier due to increased spatial reasoning and coordination. A similar measure used to track infant's growth is the development of stacking behaviors. Chimpanzees and humans can efficiently stack objects on each other by three years of age. Doing so requires knowledge of the affordance of the objects being stacked to both support objects above it and to be steady on top of the object below it [7][5]. The robot used in the present research has the developmental capacity of a child of about two years of age. These studies show that it should be possible for this robot to learn to interact with constrained space. This is a big step in this robot's natural progression towards intelligence.

III. Methodology

A high-level overview of the proposed framework for exploring constrained and traversable regions in proprioceptive space is shown in Figure 1. The robot learns from past interactions to drive future interactions by concentrating on unconstrained movements.

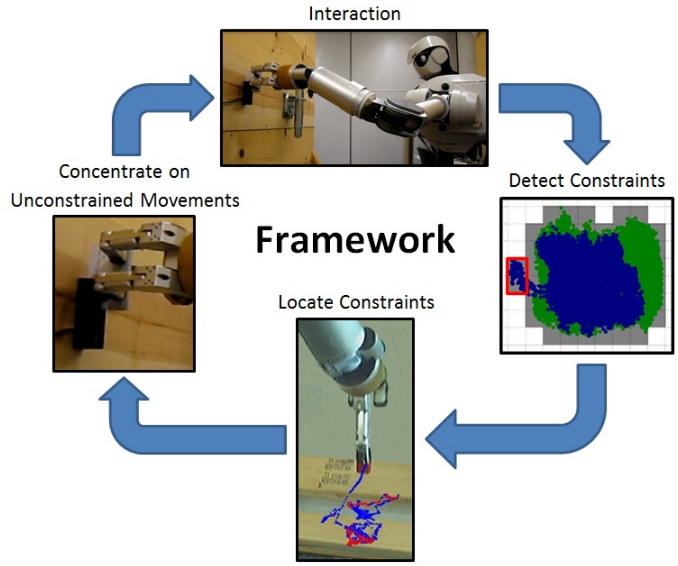

Fig. 1. The high-level structure of the framework for detecting and traversing spatial regions with constraints.

The robot performs exploratory behaviors to detect where and how its movements are constrained in space and where they are unconstrained. Proprioceptive feedback from the behaviors allows the robot to learn about unconstrained directions and favor these directions in selecting its future movements. This is beneficial because unconstrained movements are more likely to generate a desired outcome than constrained movements in a rigid space.

A. Interaction

The methods used in these sections were inspired by the field of Motor Synergies[10]. In this field, the dimensionality of data is reduced to create structures called manifolds. Manifolds contain all possible combinations of environmental variables to produce the same outcome. For example, a manifold may contain all the possible joint angle combinations possible to achieve the same position of the index finger in space. In contrast, the profile in this work represents a subset of the possible unconstrained joint torques during unconstrained movement in a given space. We believe that when the robot meets a constraint, the joint torques in the robot's arm exceed the joint torques encountered during unconstrained movement, and a comparison between the unconstrained profile and the current exploratory profile can be used to detect the state of the robot's movement.

Our method uses Principal Component Analysis (PCA). PCA is a dimensionality reducing algorithm which takes a set of N multi-dimensional vectors and returns N vectors (called principal components or PCs) ordered by the amount of variance and co-variance of the points in the data. This

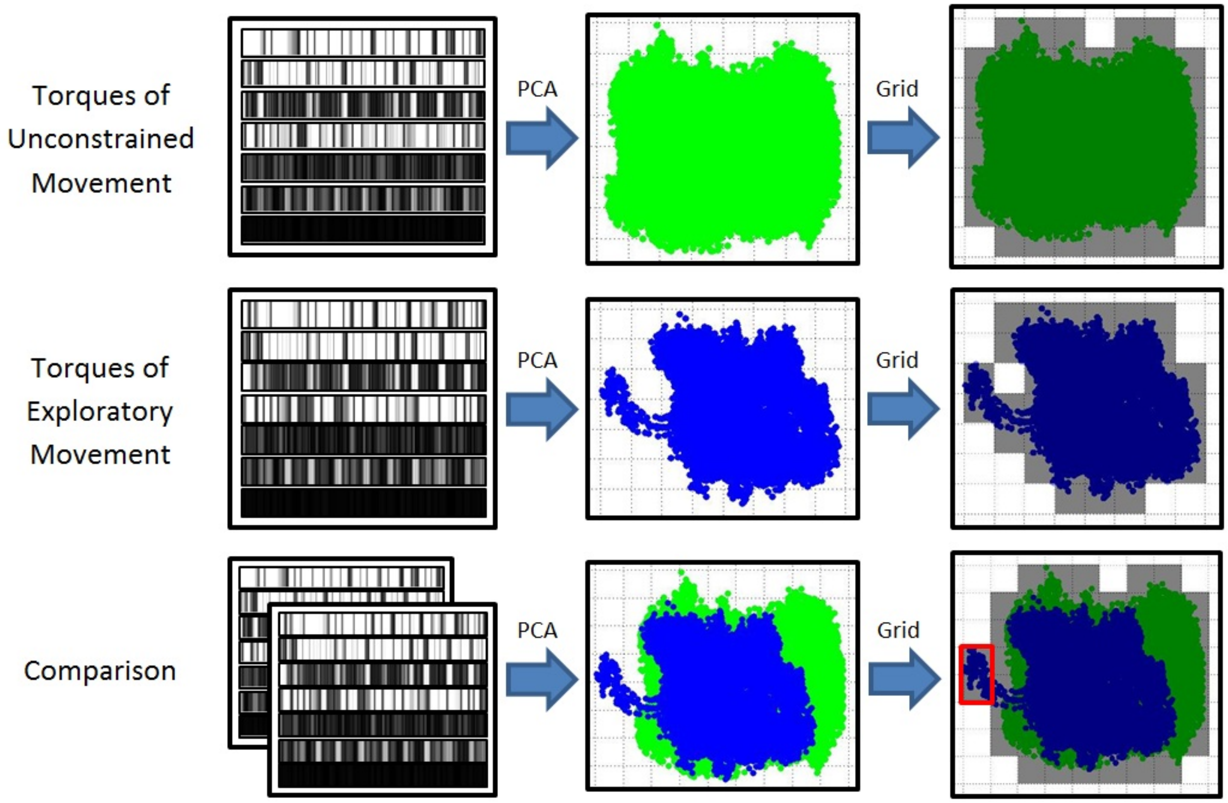

Fig. 2. The three step process to detection of constraints. The unconstrained profile is collected (top), and then the exploratory profile is collected (middle). The two profiles are compared (bottom). If the exploratory profile is contained within the unconstrained profile, it is considered similar or unconstrained. If a part of the exploratory profile is not contained within the unconstrained profile, it is considered different or constrained. In this example, the exploratory profile was constrained, as indicated by the differences in the girds, outlined in red.

scheme allows us to look at a multi-dimensional space in a representation which has fewer dimensions, while preserving the most important features of the data.

In this experiment, the robot takes a set of seven joint torque vectors, and PCA is performed on these vectors. The output is the seven principal component vectors of the data. We take the first two principal components and create the unconstrained profile. In prior research [16][20], it was found that the vast majority of variance between different hand positions (>80%) could be accounted for by just two principal components. This finding held true for our data, so we were able to plot the variance in joint torque data in two dimensions by plotting the first two principal components. Essentially, by analyzing the differences between the first two principal components, we can reduce the dimensionality of our data from seven to two. According to Flanders [3], one way to integrate perception and motor movements is to compare memories of positions of the joints. This is why we compare the PCA profiles created by exploration in constrained space to those created by exploration in unconstrained space; the profiles represent the robots memory of what was happening to its hand at specific locations. This first part of the algorithm can be visualized in the first two columns of the top row of Figure 2.

To simplify the representation of the unconstrained profile even further, the space is divided into a grid. For every cell in the grid, if any point of the unconstrained profile falls within that cell, the cell is considered a part of the unconstrained profile. The result of this process is a grid structure where the shaded cells represent the unconstrained profile, as shown in the right column of the top row of Figure 2.

B. *Constraint Detection*

The robot uses the unconstrained profile to classify its movements as either unconstrained or constrained. To do this, the robot collects all the torque readings for a particular movement. PCA is conducted on this set of torques, which creates a profile for the constrained torques. A grid is obtained in the same way as the unconstrained profile grid. This process is described in the second column of Figure 2. The two grids can now be compared as can be seen in the bottom row of Figure 2. If the exploratory profile is contained within the unconstrained profile, they are considered similar. Otherwise, the two profiles are considered different.

IV. Experimental Setup

Two experiments were performed with the robot. The card reader and magnetic strip card shown in Figure 3 were used. The card reader was mounted on a vertical board in front

Fig. 3. Magnetic card and reader.

of the robot with the card slot in a vertical orientation. The card reader was connected to a computer so that successful slides could be recorded. A slide was considered successful only if all the information on the magnetic strip of the card was read. To prevent the card from slipping in the robot's fingers a layer of tape with the adhesive exposed was placed on the top of the card and tape was placed on the tips of the robot's fingers with the adhesive not exposed as can be seen in Figure 4.

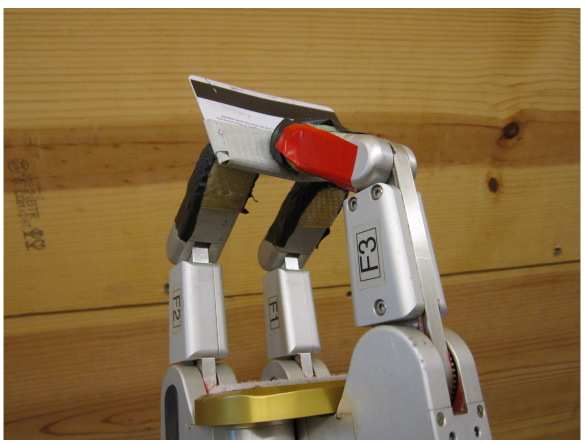

Fig. 4. Tape used to enhance the robot's grasp on the card.

For all experiments, the robot's exploration space was constrained to a triangular plane with the top point being the top of the card reader, and the two bottom points being on either side of the bottom of the card reader, as shown in Figure 5. R random joint positions were generated from the top point to a point between the bottom two points. In experiment 1, a random joint position from R was chosen and a movement to that position was performed. In Experiment 2, random points from R were chosen for some trials and a strategy, seen below, was used for others. The robot attempted performing these movements, returning to the top after each one. Before each experiment, an unconstrained profile (as described in the methodology section) was collected. The profile consisted of 100 random movements in unconstrained space. After the unconstrained profile was collected, the robot's hand, holding the card, was placed so that the card was in the card reader, close to the top point of the triangle. The robot performed random movements and data for two experiments was collected.

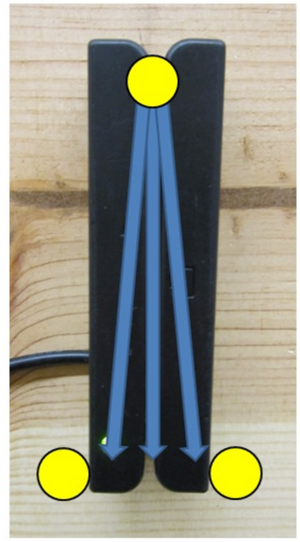

Fig. 5. Plane of exploration for the card reader defined by yellow circles. Possilbe movements indicated by blue arrows.

A. *Experiment 1*

The purpose of this experiment was to identify whether a correlation between unconstrained movements and successful slides of the card exists. The robot performed 200 random movements in the space. For each movement, the state (constrained or unconstrained) was recorded, and the card reader feedback (read or fail) was also recorded.

B. *Experiment 2*

This experiment introduced two strategies for exploration. The first was based on the constraint detection algorithm. If the robot detected a constraint during a movement, it would stop the movement, return back to the origin, and proceed onto the next movement. The second strategy was a learning scheme which favored unconstrained movements. It used the KNN algorithm with K = 3 to find the probability of unconstrained movement for each candidate joint position. The candidate with the highest probability of unconstrained movement was chosen. The robot performed 100 movements for each combination of strategies N = 4.

V. Results

The constraint detection method was applied to the data set obtained by Koonce at al.[8]. This data set contains proprioceptive data for one hundred and eighty different insertion attempts. Each attempt is categorized as successful insertion or unsuccessful insertion. Trials with successful insertions were treated as having constrained movement, and trials with unsuccessful insertions were treated as having no constrained movement. There were three insertion objects: a circular, a cross-shaped, and a hexagonal block. Each object had a corresponding hole, and would fit only in that hole. Twenty insertion attempts were performed for each object-hole combination. Two unconstrained trials were chosen at random to build the unconstrained profile, and

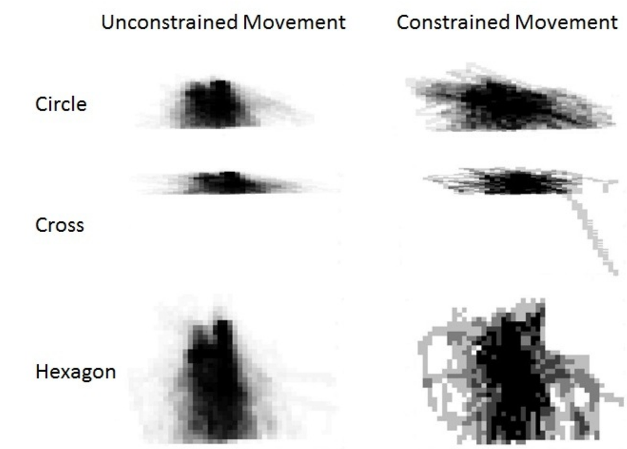

Fig. 6. Comparison of constrained and unconstrained profiles.

PCA was performed on every other trial using the principle components of the two unconstrained trials. The data set was separated into three categories according to the shape of the insertion object (circle, cross, or hexagon), and each category was separated into two subcategories (unconstrained and constrained). The obtained results are displayed in Figure 6. In all three cases, differences between the constrained and unconstrained profiles can be observed. Furthermore, the pattern appears to be generalizable, as the unconstrained profiles are more centered, while the constrained profiles exceed the boundaries of the unconstrained profiles. This is consistent with this work's proposal that the unconstrained profile represents the space of most joint torque combinations during unconstrained movement and the introduction of a constraint results in joint torque combinations outside the unconstrained space.

A distance matrix between all trials was also created using our method. Every trial was treated as unconstrained and an unconstrained profile from that trial was created. Then, all trials were treated as exploratory, and the number of elements that were outside the unconstrained profile was found. This difference would then be recorded with the row being the number of the unconstrained trial and the column being the number of the exploratory trial. The distance matrix for all trials with the circular block can be seen in Figure 7. The trials were grouped by category, with the constrained trials being the lower numbered trials and the unconstrained trials being the higher numbered trials. White represents high similarity and darker shades represent lower similarity. The constrained trials are all darker shades meaning they were very dissimilar from all other trials, and that our method algorithm was able to separate the constrained trials from the unconstrained trials. Distance matrices were also constructed for the cross shaped and hexagonal blocks. However, since the experiment focused on insertions and not constraints, the data only indicates when objects were inserted into holes. These two shapes could become partially inserted and therefore constrained but be recorded as unconstrained because the block was not fully inserted. Therefore the results from that analysis are inconclusive and will not be included.

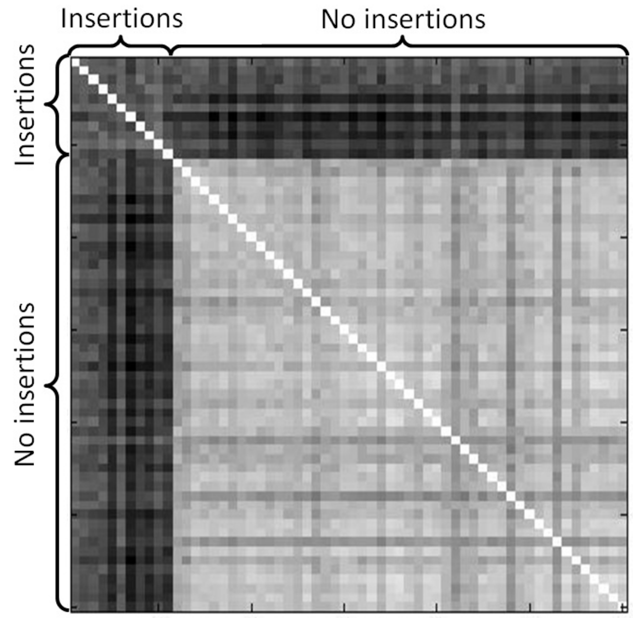

Fig. 7. Distance matrix for sixty trials with the circle shaped block. The first eleven trials are constrained and the rest are unconstrained.

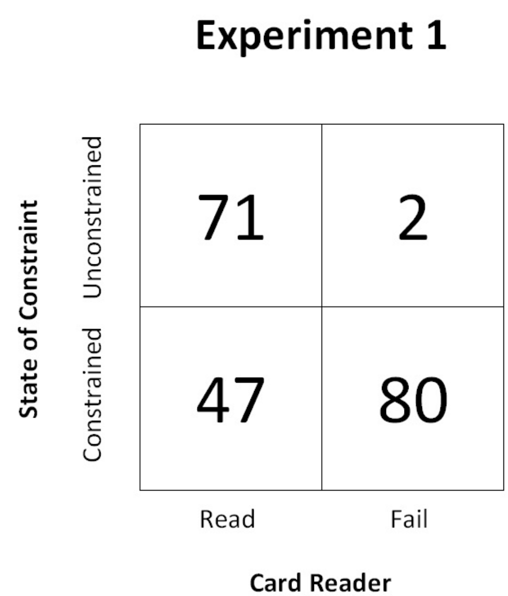

Fig. 8. Results from Experiment 1. 200 random movements were performed. 97% of all unconstrained movements successfully read the card and only 3% failed.

A. Experiment 1

The results for this experiment can be seen in Figure 8. The results of this experiment show that unconstrained movement can be associated with successful scanning of the card. Out of the 200 trials, 73 unconstrained movements were performed. During these movements, the card was successfully scanned 71 times and was unsuccessfully scanned 2 times. There appears to be a correlation between unconstrained movements and sliding the card as 97% of unconstrained movements resulted in a successful read.

Experiment 2

Learning	Constraint Detection	Successful Slides (out of 100 movements)
On	On	95
On	Off	75
Off	On	10
Off	Off	*

Fig. 9. Results from Experiment 2. *The robot failed to complete the entire trial in 10 out of 10 trials when learning and constraint detection were turned off; on average, the robot attempted to slide the card 12.7 times and succeeded in doing so 4.5 times before it removed the card from the reader.

B. Experiment 2

The results of this experiment can be seen in Figure 9. A relationship between the learning strategy, the constraint detection strategy, and successful scanning of the card can be observed. Each trial of this experiment consisted of 100 movements. When the learning strategy and the detection strategy were both used, the robot performed 95 successful slides. When the learning strategy was on and the constraint detection was off, the robot performed 75 successful slides. When the learning strategy was turned off and the constrained detection strategy was on, the robot performed 10 successful slides.

When both strategies were turned off, the robot was not able to complete a trial as it exerted too much force and bent the card out of the card reader. Ten attempts to complete this trial were performed and the robot failed all ten at an average of 53.27 seconds after the beginning of the experiment. On average, the robot slid the card 12.7 times and was succesful 4.5 times. These attempts are similar to the trials in Experiment 1, however, the speed of movements of the robot was increased in this experiment to increase efficiency. Because of this, the force exerted during constrained movements was greater and the robot forced the card out of the card reader. These results show that the learning strategies had a significant effect on the robot's ability to perform the task. Without the learning strategies, the speed increase made the task almost impossible to complete for the robot. With the strategies, the robot was able to complete a significantly higher number of succesful slides.

VI. Conclusion

This work proposed a constraint detection algorithm based exclusively on proprioception. The algorithm was tested on a previous data set with promising results. It was also tested on a card reader. The results of those experiments indicate that the algorithm can be used to associate unconstrained movement with task completion. Furthermore, learning strategies were developed which increased the robot's ability to complete the task succesfully. Our algorithm can be used to enable robots to interact with environments in conditions where visual models cannot be created.

VII. Future Work

The method we have presented in this paper can be applied in many different ways. We have shown that it can be used to explore constraints and learn from them. We foresee this method being used in other areas as well. One possible extension of this method is detecting that a key has been inserted and finding out which way it turns. Using the exploration featured in our method, the robot would be able to operate a lock. Furthermore, this method could allow robots to associate different keys with different locks. Observation of how the key turns in the lock would allow the robot to determine which key works in which lock.

Spatial semantic hierarchies could be constructed using our method. For example, opening a door allows one to open the door and access what lies behind it. Our method could detect that, when closed, the door acts as a constraint, blocking the robot from getting behind it. Once open, the door no longer acts as a constraint in this manner. A hierarchy could be built showing that opening the door allows the robot to access what is behind it based on the removal of the door's constraint.

Mobile robots could be able to use this method for assistance in navigation in multiple ways. Robots needing to get through doors could use our method to learn how to use door handles and how doors open. Exploration would allow the robot to open a door regardless of whether the door opens in or out. The robot could also learn to manipulate many different door handles by exploring.

This work was funded in part by NSF Research Experience for Undergraduates (REU) Grant IIS-0851976.

References

[1] H. Bruyninckx, S. Dutre, and J. De Schutter. Peg-on-hole: a model based solution to peg and hole alignment. In *Robotics and Automation, 1995. Proceedings., 1995 IEEE International Conference on*, volume 2, pages 1919 –1924 vol.2, may 1995.

[2] K. Connolly and M. Dalgleish. The emergence of a tool-using skill in infancy. *Developmental Psychology*, 25(6):894 – 912, 1989.

[3] M. Flanders. What is the biological basis of sensorimotor integration? *Biological Cybernetics*, 104(1-2):1–8, 2011.

[4] E. Gibson. Exploratory behavior in the development of perceiving, acting, and the acquiring of knowledge. *Annual Review of Psychology*, 39(1):1–41, 1988.

[5] M. Hayashi. Stacking of blocks by chimpanzees: developmental processes and physical understanding. *Animal Cognition*, 10:89–103, 2007.

[6] M. Hayashi and T. Matsuzawa. Cognitive development in object manipulation by infant chimpanzees. *Animal Cognition*, 6:225–233, 2003.

[7] M. Hayashi and H. Takeshita. Stacking of irregularly shaped blocks in chimpanzees (*Pan troglodytes*) and young humans (*Homo sapiens*). *Animal Cognition*, 12:49–58, 2009.

[8] P. Koonce, V. Dutell, J. Farrington, V. Sukhoy, and A. Stoytchev. Toward learning to solve insertion tasks: A developmental approach using exploratory behaviors and proprioception. In *In Proceedings of AAAI 2011*, 2011.

[9] B. Koslowski and J. Bruner. Learning to use a lever. *Child Development*, 43(3):pp. 790–799, 1972.

[10] M.L. Latash. *Neurophysiological basis of movement*. Neurophysiological Basis of Movement. Human Kinetics.

[11] B. Mayton, L. LeGrand, and J. Smith. Robot, feed thyself: Plugging in to unmodified electrical outlets by sensing emitted ac electric fields. In *International Conference on Robotics and Automation*, pages 715–722, 2010.

[12] W. Meeussen, M. Wise, and S. et al. Glaser. Autonomous door opening and plugging in with a personal robot. In *ICRA*, 2010.

[13] J. Morrow, B. Nelson, and P. Khosla. Vision and force driven sensorimotor primitives for robotic assembly skills. *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, 3:3234, 1995.

[14] H. Ornkloo and C. von Hofsten. Fitting objects into holes: On the development of spatial cognition skills. *Developmental Psychology*, 43(2):404–416, 2007.

[15] W. Paetsch and G. von Wichert. Solving insertion tasks with a multifingered gripper by fumbling. In *Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on*, pages 173 –179 vol.3, may 1993.

[16] M. Santello, M. Flanders, and J. Soechting. Postural hand synergies for tool use. *Journal of Neuroscience*, 18(23):10105–10115, 1998.

[17] K. Shutts, H. Ornkloo, C. Von Hofsten, R. Keen, and E. Spelke. Young children's representations of spatial and functional relations between objects. *Child Development*, 80(6):1612–1627, 2009.

[18] R. Suarez, L. Basanez, and J. Rosell. Using configuration and force sensing in assembly task planning and execution. *Assembly and Task Planning, IEEE International Symposium on*, 0:0273, 1995.

[19] L. Van Leeuwen, A. Smitsman, and C. Van Leeuwen. Affordances, perceptual complexity, and the development of tool use. *Journal of Experimental Psychology: Human Perception and Performance*, 20(1):174–191, 1994.

[20] E. Weiss and M. Flanders. Muscular and Postural Synergies of the Human Hand. *J Neurophysiol*, 92(1):523–535, July 2004.