Which Object Comes Next? Grounded Order Completion by a Humanoid Robot

C. Schenck, J. Sinapov, A. Stoytchev

Developmental Robotics Laboratory Iowa State University, Ames IA, USA

Emails: cschenck@iastate.edu jsinapov@iastate.edu alexs@iastate.edu

Abstract: *This paper describes a framework that a robot can use to complete the ordering of a set of objects. Given two sets of objects, an ordered set and an unordered set, the robot's task is to select one object from the unordered set that best completes the ordering in the ordered set. In our experiments, the robot interacted with each object using a set of exploratory behaviors, while recording feedback from two sensory modalities (audio and proprioception). For each behavior and modality combination, the robot used the feedback sequence to estimate the perceptual distance for every pair of objects. The estimated object distance features were subsequently used to solve ordering tasks. The framework was tested on object completion tasks in which the objects varied by weight, compliance, and height. The robot was able to solve all of these tasks with a high degree of accuracy.*

Keywords: *Developmental robotics, object exploration, grounding.*

1. Introduction

Humans can detect order in an unordered set of objects at a very early age. Ordering tasks frequently appear on modern intelligence tests [7, 8]. They are also tightly integrated in many educational methodologies. For example, in the Montessori method [12], a 100-year-old method of schooling for children that has been shown to outperform standard methods [10, 11], children are encouraged to solve different object ordering tasks with specialized toys [16]. These strongly suggest that the ability to discover orderings among sets of objects is an important skill. Indeed, studies in

psychology have revealed that this skill is learned at a very early age [22, 6, 2, 3].

Because order completion skills are so important for humans they should be important for robots that operate in human environments as well. Previous research has shown that robots can successfully form object categories [15, 14, 13, 24] and solve the odd-one-out task [19]. Object ordering tasks, however, have not received a lot of attention from the robotics community to date.

This paper proposes a method for discovering orderings among groups of objects. The experiments were conducted with an upper-torso humanoid robot, which interacted with the set of objects using a set of stereotyped exploratory behaviors. The robot recorded both auditory and proprioceptive data during each interaction and then extracted features from the sensory records. Using the extracted features for each object, the robot was able to estimate a pairwise distance matrix between every pair of objects. Then given three objects that form an ordered set, the robot’s model was queried to pick one object from another group of four to complete the ordering in the first set. The results show that the robot was able to pick the correct object that completes the ordering with a high degree of accuracy and that different exploratory behaviors and sensory modalities are required to capture different ordering concepts.

2. Related work

Object ordering tasks appear on multiple intelligence tests. For example, on the Intelligence and Development Scales (IDS) test [7], children are asked to sort lines of varying length. In a more common test, the Wechsler Intelligence Scale for Children (WISC) [8], participants are asked to place images from a story into a logical sequence. While it is not currently feasible for a robot to understand the events taking place in an image, these two tests show that, given an understanding of the objects, knowledge of how to order them is a strong indicator of intelligence.

Several studies have shown that young children have a fundamental understanding of the concepts underlying ordering. *Graham et al.* [6] found that children between the ages of 2 and $4\frac{1}{2}$ can easily judge an object as “big” or “small” when compared to another object. Two studies by *Ebeling* and *Gelman* [2, 3] found similar results. Interestingly, all three studies found that children were much better able to judge an object as “big” or “small” when compared with immediately viewable objects as opposed to making the judgment based on the object’s absolute size [6], its normative size (i.e., how big it is compared to the typical object in the category) [2], or its functional size (i.e., how big it is in relation to the function it is to perform) [3]. The ability to compare an object to other directly viewable objects is a prerequisite for successfully performing the task of ordering and since it is present more strongly than other types of comparisons (absolute, normative, or functional) at such an early age, it must be fundamental to intelligence.

In another study with 1 to 3-year-olds, *Sugrman* [22] found that the order in which children interact with objects tends to be influenced by the class and perceptual similarity of the current object to the previously explored object. Additionally, it was observed that the older children relied less on the class of the object to pick the

Fig. 1. The three sets of objects used in the experiments

next object and more on perceptual similarity. This paper uses a similar method to determine orderings. The perceptual distances between each pair of objects is used to determine the best object to complete the ordering.

In machine learning, the problem of ranking (i.e., placing a set of data in the correct order) has been well studied [1]. There are many algorithms that can solve ranking with a high degree of accuracy. It is difficult to use standard ranking methods, however, to perform order completion tasks, especially when the number of objects is small. Additionally, standard ranking methods are often supervised or semi-supervised. On the other hand, the method proposed in this paper solves the task in both unsupervised and supervised settings.

Not a lot of research has been done in robotics on discovering orderings in small groups of objects. Measuring the similarity between objects, however, is a common way to solve tasks in robotics. There have been numerous experiments that have demonstrated robots' ability to measure perceptual as well as functional object similarity for varying tasks [15, 14, 13, 24, 23, 18, 4, 19]. Multiple studies [13, 17, 20, 18] have used the similarity of perceptual features to categorize objects in an unsupervised manner. In [19], perceptual distances between objects were used to solve the odd-one-out task. This paper builds on this previous work by proposing a method to solve the order completion task.

3. Experimental platform

All experiments were performed with the lab's upper-torso humanoid robot, which has two 7-DOF Barrett Whole Arm Manipulators (WAMs) as its actuators, each with an attached Barrett Hand. The robot captured proprioceptive information from the built-in sensors in the WAM that measure the angles and the torques applied to each joint at 500 Hz. The robot also captured audio data through an Audio-Technica U853AW cardioid microphone mounted in its head at the standard 16-bit/44.1 kHz over a single channel.

The robot was tested on three ordering concepts: ordering by *weight*, ordering by *compliance*, and ordering by *height*. Fig. 1 shows the three sets of objects that were used in the experiments. The first two are standard Montessori toys. The *weight cylinders* are composed of six pairs of objects (for a total of twelve objects) that vary by weight, with the objects in each pair having the same weight. All the *weight cylinders* are functionally identical except for their weight. The *pressure cylinders* are composed in a similar manner (six pairs of objects) except that they vary by the

amount of pressure required to depress the rod on top of the object. The *cones* and *noodles* are composed of five green, styrofoam cones of varying sizes and five pink, foam pieces (cut from a water noodle) ranging in size from small to large. Because the object's in the first two sets are visually identical, this task cannot be solved with vision alone. In fact, the robot did not use vision at all to solve the ordering task.

The robot performed nine behaviors on each of the objects: *grasp*, *lift*, *hold*, *shake*, *drop*, *tap*, *poke*, *push*, and *press*. Additionally, the behavior *rattle* was performed on the *weight cylinders* and the *pressure cylinders*. Each behavior was encoded as a trajectory in joint-space for the left arm using the Barrett WAM API and executed using the default PID controller. All behaviors were performed identically on each object with the exception of *grasp* and *tap*, which were adjusted automatically based on the current visually detected location of the object. Fig. 2 shows the robot performing each behavior on one of the pressure cylinders.

At the start of each trial, the experimenter placed one of the objects on the table in front of the robot. The robot then performed the exploratory behaviors on the object, with the experimenter placing the object back on the table if it fell off. This was repeated five times for each of the *cones* and *noodles* and ten times for the rest of the objects. The data for the *cones* and *noodles* was collected at an earlier time than for the rest of the objects, which is why only five repetitions were done and the behavior *rattle* was not performed on them. During each behavior, the robot recorded proprioceptive data in the form of joint torques applied to the arm over time and auditory data in the form of a wave file. Visual input was used only to determine the location of the object for the *grasp* and *tap* behaviors.

Fig. 2. The ten exploratory behaviors that the robot performed on the objects. From left to right and top to bottom: *grasp*, *lift*, *hold*, *shake*, *drop*, *tap*, *poke*, *push*, *press*, and *rattle*. The *rattle* behavior wasn't performed on the *cones* and *noodles*. The object in this figure is one of the pressure cylinders. After some of the behaviors (e.g., *drop*), the object was moved back to the red marker location on the table by the experimenter.

4. Feature extraction

4.1. Sensorimotor feature extraction

The auditory feedback from each behavior was represented as the Discrete Fourier Transform (DFT) of the sound’s waveform, computed using 33 frequency bins. Thus, each interaction produced a $33 \times n$ matrix, where each column represented the intensities for different frequencies at a given point in time (i.e., n was the number of samples). The DFT matrix was further discretized uniformly into 10 temporal bins and 10 frequency bins. Thus, the auditory feature vector for each interaction was a $10 \times 10 = 100$ dimensional real-valued vector.

The proprioceptive feedback was represented as 7 time series of detected joint-torques, one for each of the robot’s joints. To reduce the dimensionality of the data, each of the series was uniformly discretized into 10 temporal bins. Thus, the proprioceptive features for each interaction were represented by a $7 \times 10 = 70$ dimensional real-valued vector. As described next, the computed auditory and proprioceptive features were used to estimate the pairwise distances for each pair of objects.

4.2. Object feature extraction

Let $\mathcal{C}$ be the set of sensorimotor contexts, i.e., each $c \in \mathcal{C}$ corresponds to a behavior-modality combination (e.g., audio-shake), and let $\mathcal{O}$ denote the full set of objects. The goal of the object feature extraction routine is to compute a distance matrix $\mathbf{W}^c$ such that each entry $W_{ij}^c \in \mathbb{R}$ encodes how perceptually different objects o_i and o_j are in sensorimotor context c . Let the set $\mathcal{X}_i^c = [x_1, \dots, x_D]_i^c$ contain the sensorimotor feature vectors detected for each of the D exploratory trials with object o_i in context c . The distance between two objects o_i and o_j in context c can be represented by the expected distance between the feature vectors in $\mathcal{X}_i^c$ and the feature vectors in $\mathcal{X}_j^c$, i.e.,

$$W_{ij}^c = \mathbf{E}[d_{L2}(x_a, x_b) | x_a \in \mathcal{X}_i^c, x_b \in \mathcal{X}_j^c],$$

where d_{L2} is the L2-norm distance function. This expectation is estimated by:

$$W_{ij}^c = \frac{1}{|\mathcal{X}_i^c| \times |\mathcal{X}_j^c|} \sum_{x_a \in \mathcal{X}_i^c} \sum_{x_b \in \mathcal{X}_j^c} d_{L2}(x_a, x_b).$$

The result is a set $\mathcal{W}$ of object distance matrices, where each $\mathbf{W}^c \in \mathcal{W}$ encodes the pairwise perceptual distance for each pair of objects in $\mathcal{O}$. The next section describes how these matrices can be used to decide which one of a given set of objects best completes a given order.

5. Methodology

5.1. Problem formulation

Each order completion task is formulated as follows. Let $\mathcal{O}$ denote the set of objects explored by the robot. Let $\mathcal{L}$ denote an *ordered subset* of $\mathcal{O}$, i.e., $\mathcal{L} = o_1, o_2, \dots, o_N$ where each $o_i \in \mathcal{O}$. Furthermore, let $\mathcal{G} \subset \mathcal{O}$ be an unordered set of M objects denoting the set of candidate objects that could be selected to complete the order. Finally, let

$\mathcal{W}$ be a set of distance matrices such that for a given sensorimotor context c , the $|O| \times |O|$ matrix $\mathbf{W}^c \in \mathcal{W}$ encodes the pairwise object distances in that context.

In this setting, the task of the robot’s model is to select one object from $\mathcal{G}$ that correctly completes the order specified by the ordered set $\mathcal{L}$. The idea behind the approach presented here is to define an objective function that can evaluate the quality of a proposed order and use that function to select an object from the set $\mathcal{G}$. The next sub-section describes the objective function as well as how that function is used to pick an object that completes the order.

5.2. Selecting the best order completion candidate

Let $q(\mathcal{L}, \mathbf{W}^c)$ denote the objective function that measures the quality of the order $\mathcal{L}$ with respect to the matrix $\mathbf{W}^c$. That function is defined as:

$$q(\mathcal{L}, \mathbf{W}^c) = \sum_{o_i \in \mathcal{L}} \sum_{o_j \in \mathcal{L}} (W_{ij}^c - d(o_i, o_j, \mathcal{L}))^2,$$

where the function d is defined as

$$d(o_i, o_j, \mathcal{L}) = \sum_{r=o_i \dots o_{(j-1)} \in \mathcal{L}} W_{r(r+1)}^c.$$

In other words, the function d approximates the distance between objects o_i and o_j by summing up the distances between adjacent elements in the ordered set $\mathcal{L}$. Thus, the function q measures the squared difference between the true distance matrix and the one approximated by the proposed ordering. It is used by the robot’s model to complete a given ordered set of objects as follows. For each object o_k from the unordered set $\mathcal{G}$, let $\{\mathcal{L}, o_k\}$ denote the *ordered set* of objects produced by adding object o_k to the end of the ordered set $\mathcal{L}$. In this setting, the model selects the object o_k that maximizes the objective function $q(\{\mathcal{L}, o_k\}, \mathbf{W}^c)$.

5.3. Order completion using multiple sensorimotor contexts

The method presented so far can only use one distance matrix $\mathbf{W}^c$ that is specific to one sensorimotor context c . For many tasks, however, it may be desirable to use multiple sources of information about how objects relate to each other. For example, if the given ordered set of objects $\mathcal{L}$ is ordered by weight, there may be several exploratory behaviors that capture relevant proprioceptive information for solving the task (e.g., *lifting* and *holding* in place).

The set $\mathcal{W}$ contains multiple matrices encoding the pairwise object dissimilarities computed for a given set of sensorimotor contexts. For each object $o_k \in \mathcal{G}$, let the function $\text{completes}(\mathcal{L}, o_k, \mathbf{W}^c)$ return 1 if o_k is selected as the object completing the order and 0 otherwise. Given the set of all matrices $\mathcal{W}$, the ordered set $\mathcal{L}$, and the candidate set $\mathcal{G}$, the model selects the object $o_k \in \mathcal{G}$ that maximizes the following function:

$$\text{score}(o_k) = \sum_{\mathbf{W}^c \in \mathcal{W}} w_c \times \text{completes}(\mathcal{L}, o_k, \mathbf{W}^c),$$

where w_c is a weight that encodes the relevance of sensorimotor context c .

In the experiments described in the next section, three weighting methods are evaluated. Whereas everything in this paper so far has been unsupervised, two of

these weighting methods are supervised (methods 2 and 3). In the first method, the weights are *uniform*. In other words, for all c , $w_c = 1.0$.

In the second method, the weights are set to the estimated accuracy of using sensorimotor context c to solve the specific ordering task. In other words, the robot's model estimates the accuracy of a context c by running the method described in the previous subsection on a training set of tasks of the form $[\mathcal{L}, \mathcal{G}]$ for which the correct answers are known. Once the weights for all contexts have been estimated, the model uses those weights on subsequent tasks for which the answers are not known in advance.

The third method that was used to combine sensorimotor contexts is boosting. It was implemented using the AdaBoost algorithm [5]. It is briefly summarized here. Given a set of m tasks $[\mathcal{L}_1, \mathcal{G}_1], [\mathcal{L}_2, \mathcal{G}_2], \dots, [\mathcal{L}_m, \mathcal{G}_m]$ for which the correct answers $o_k^1 \in \mathcal{G}_1, o_k^2 \in \mathcal{G}_2, \dots, o_k^m \in \mathcal{G}_m$ are known, initialize the training weights as $D_1(i) = \frac{1}{m}$ for $i = 1, \dots, m$. For each iteration $t = 1, \dots, T$, select the sensorimotor context $c^*(t)$ such that $c^*(t) = \arg \min_{c \in C} \xi_c$. The error ξ_c of a context c is computed as

$$\xi_c = \sum_{i=1}^m D_t(i) [1 - \text{completes}(\mathcal{L}_i, o_k^i, \mathbf{W}^c)],$$

where $o_k^i \in \mathcal{G}_i$ is the object that correctly completes the ordering $\mathcal{L}_i$. Next, the parameter α_t is computed as a function of $\xi_{c^*(t)}$ as follows

$$\alpha_t = \frac{1}{2} \ln \frac{1 - \xi_{c^*(t)}}{\xi_{c^*(t)}},$$

where $\xi_{c^*(t)}$ is the error of the selected context in iteration t . After each iteration, the training weights for all $i = 1, \dots, m$ are updated as follows

$$D_{t+1}(i) = D_t(i) \exp \left[-\alpha_t (2 * \text{completes}(\mathcal{L}_i, o_k^i, \mathbf{W}^{c^*(t)}) - 1) \right],$$

where $\mathbf{W}^{c^*(t)}$ is the object distance matrix of the context selected during iteration t , and then normalized such that they sum to 1. It is worthwhile to note that the expression $-\alpha_t (2 * \text{completes}(\mathcal{L}_i, o_k^i, \mathbf{W}^{c^*(t)}) - 1)$ comes out to $+\alpha_t$ if context $c^*(t)$ incorrectly predicts the object to complete the ordering $\mathcal{L}_i$ and $-\alpha_t$ otherwise. In essence, the training weights are altered such that tasks that context $c^*(t)$ is incorrect on are weighted higher and tasks that it is correct on are weighted lower.

Finally, the weight w_c for each sensorimotor context is computed by

$$w_c = \sum_{t=1}^T \alpha_t [c \equiv c^*(t)],$$

where $[c \equiv c^*(t)]$ is 1 if c was chosen during iteration t and 0 otherwise. In the experiments described in this paper, T was set to 50. Results did not change significantly with a higher value for T .

5.4. Evaluation

The model was evaluated independently on each of the three ordering concepts. Fifty tasks were randomly sampled for each concept as follows: four objects were sampled from the set O such that there existed a clear ordering amongst them (e.g., for the

weight cylinders, two objects from the same pair would not be sampled together). The objects were then ordered (with the direction, forward or backward, determined randomly) and the last object was removed. Thus, the first three ordered objects formed the ordered set $\mathcal{L}$ for the given task. Three more objects were randomly sampled from the remaining objects in $\mathcal{O}$ such that none validly completed the ordering. These three objects, combined with the removed object, formed the set $\mathcal{G}$. The performance of the robot’s model was evaluated in terms of accuracy, i.e., the number of tasks for which the robot’s model picked the correct object divided by the total number of tasks.

For each concept, the performance of each sensorimotor context was evaluated. The accuracy was also computed as more and more contexts were used by the model. To estimate the context weights and to train the boosting method, five-fold cross-validation was performed with the 50 sampled tasks (i.e., 10 tasks were randomly assigned to each fold). The model was also evaluated as the number of tasks used for training was varied from 1 to 49. In this case, all contexts were used.

6. Results

6.1. An example order completion task

Fig. 3 shows an example task in which the robot’s model is tasked with completing an order of three objects that are ordered by height. In this case, the ordered input set, $\mathcal{L}$, consists of three pink noodles, while the candidate set, $\mathcal{G}$, contains four objects – three cones and one noodle, such that only one of them is taller than the last element in $\mathcal{L}$. In this specific case, the input distance matrix encoded the perceptual similarity of the objects in the *press-proprioception* sensorimotor context. The figure shows an ISOMAP [25] embedding of the distance matrix, which makes it easy to see that the matrix encodes an order between the objects. For this task, the model correctly picked the cone from the set $\mathcal{G}$ that is taller than the tallest noodle object in $\mathcal{L}$. The next subsection describes a quantitative evaluation of the model in which each sensorimotor context is evaluated on each of the three ordering tasks.

Fig. 3. An example task. The box on the left shows both the ordered set $\mathcal{L}$ and the unordered set of objects $\mathcal{G}$ to choose from. The plot on the right shows the ISOMAP embedding of the distance matrix between the objects. The blue circles denote the three objects in $\mathcal{L}$, the red circle denotes the object in $\mathcal{G}$ that is selected to complete the order.

6.2. Ordering objects using a single sensorimotor context

For the first experiments, the performance of the model was evaluated using a single sensorimotor context. Fig. 4 shows the accuracy for each context on each of the 3 concepts. As expected, *lift* (100%), *drop* (100%), *hold* (98.0%), *shake* (100%), and *rattle* (98.0%) for *proprioception* perform very well on the task of ordering objects by weight. This is likely because the robot was supporting the full weight of the object with its arm while performing these behaviors. For the *pressure cylinders*, *proprioception-lift* (100%) and *proprioception-tap* (98.0%) achieve high performance. The reason for this is likely due to the weight and moment of inertia differences in the objects caused by the different springs inside the *pressure cylinders*. *Proprioception-press* was able to achieve 100% accuracy on the *cones* and *noodles* task as was expected since the moment at which the arm touched the object varied depending on the object’s height. The other sensorimotor contexts did not perform as well, with *proprioception-push* (84.0%) being the next highest performing context.

Fig. 4. The accuracy of each context for each of the 3 concepts. Darker values indicate lower accuracies with solid black being 0%; lighter values indicate higher accuracies with solid white being 100%.

6.3. Ordering objects using multiple sensorimotor contexts

Fig. 5 shows the performance of the robot on each concept as the number of contexts varies from 1 to $|C|$ when using the uniform, weighted, and boosted combination methods as described in section 5.3. The accuracy when picking just the single-best context (based on the training tasks) is also shown for comparison. As the number of combined contexts increases, the average accuracy also increases, which is consistent with our previous results [17]. Additionally, in every case the weighted combination method outperforms the uniform combination method. Also the boosted method always does at least as well as the weighted method and in most cases outperforms it. For the *weight cylinders* (Fig. 5a), the weighted method reaches 98.0% accuracy and the boosted method reaches 100% accuracy when all contexts are used. For the *pressure cylinders* (Fig. 5b) the weighted and boosted combination methods reach 100% when all contexts are used. For the *cones* and *noodles* (Fig. 5c), the single-best context is able to achieve 100%, but when all the sensorimotor contexts are combined the robot was only able to achieve 72.0% accuracy using the weighted method. Using the boosted method, however, it was able to reach 100%. We believe that since for the

Fig. 5. The accuracy as the number of contexts is increased. The dark-gray line is the accuracy when picking the single-best context; the black line is the accuracy when using uniform weights to combine contexts; the gray line is the accuracy when the contexts are weighted in proportion to their individual accuracies; and the light-gray line is the accuracy achieved when using AdaBoost to learn the weights.

height concept, unlike for the other two, there was only one context that performed well, the noise from combining underperforming contexts outweighed the single best performing context for the weighted method, but the boosted method was able to learn this and weight the best context higher.

Fig. 6 shows the average accuracy as the number of tasks used for training is varied from 1 to 49 when combining all sensorimotor contexts. Again the single-best context (based on the training tasks) is shown for comparison. In every case, the weighted method converges after no more than 6 training tasks are used to estimate the weights. The boosted method always achieves 90% accuracy after no more than 4 training tasks and 95% accuracy after no more than 7. For *weight*, the boosted method and the weighted method converge at approximately the same rate. For *height*, the boosted method outperforms the weighted method by a large margin. For *compliance*, the boosted method converges slower than the weighted method (weighted reaches 100% after 3 tasks are used while boosted doesn't reach 100% until 40 tasks are used). This is likely related to the result in fig. 5, where *compliance* is the only task in which the uniform combination method reaches 100%. Interestingly, while the boosted method and the single-best context (as determined by the training set)

Fig. 6. The average accuracy as the number of tasks used for training is increased. The dark-gray line is the accuracy achieved when picking just the single-best context; the gray line is the accuracy achieved when the contexts are weighted in proportion to their individual accuracies; and the light-gray line is the accuracy achieved when using boosting. The results are averaged over 50 sets of training tasks for each size from 1 to 49.

converge to 100% for all three concepts, the boosted method converges much quicker for both the *weight cylinders* and *pressure cylinders*, and at about the same rate for the *cones* and *noodles*.

7. Conclusion and future work

In this paper we presented a theoretical model for performing order completion. We evaluated this model using an upper-torso humanoid robot on three concepts: *weight*, *compliance*, and *height*. The results show that the robot was able to select objects to complete orderings with a high degree of accuracy. For each concept, there existed at least one sensorimotor context that was able to achieve 100% accuracy, and there were multiple such contexts for *weight* and *compliance*. When combining sensorimotor contexts, on average, the best performance was achieved when all contexts were used, though in every case the best single context did at least as well or better. This suggests that when completing an ordering determined predominantly by only one property (e.g., weight), if there exists at least one sensorimotor context that is able to capture that property, then its predictions will typically align with the true ordering.

Given these results, what strategy should the robot use to solve a novel order completion task? The results clearly show that the boosted combination method is the best strategy for combining sensorimotor contexts because it always performs as well as or better than every other method and because it usually takes very few training tasks to train. The methodology used in this paper builds on our previous work, in which we have shown that stereotyped exploratory behaviors can be used to detect functional similarities between tools [18], perform object recognition [17], perform object categorization [20], recognize surface textures [21], solve the odd-one-out task [19], and now solve the order completion task. These results suggest that a wide variety of tasks can be solved using a library of task-specific algorithms applied on a common set of sensorimotor data extracted from exploratory behaviors.

A limitation of the method described in this paper is that while it can solve order completion tasks in which the order is ascending or descending by one property, it cannot solve more complicated tasks from the general domain of sequence completion. Therefore, future work will need to consider methods to solve completion tasks in which the transitions between elements are more complex than simply increasing or decreasing. Future work can also build on this model and others such as [19] and [9] that analyze the structure among groups of objects by using the discovered properties to scaffold learning of more complex concepts. Pursuing this line of research could allow robots to learn more complicated concepts that can be represented in terms of simpler concepts such as the ones explored in this paper.

References

1. Chapelle, O., Y. Chang, T. Liu. Future Directions in Learning to Rank. – In: JMLR: Workshop and Conference Proceedings, Vol. **14**, 91-100.
2. Ebeling, K., S. Gelman. Coordination of Size Standards By Young Children. – Child Development, Vol. **59**, No 4, 1988, 888-896.

3. Ebeling, K., S. Gelman. Children's Use of Context in Interpreting Big And Little. – *Child Development*, Vol. **65**, No 4, 1994, 1178-1192.
4. Forssén, P., D. Meger, K. Lai, S. Helmer, J. Little, D. Lowe. Informed Visual Search: Combining Attention And Object Recognition. – In: *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, 935-942.
5. Freund, Y., R. Schapire. A Desicion-Theoretic Generalization of On-Line Learning And an Application to Boosting. – In: *Computational learning theory*, Springer, 23-37.
6. Graham, F., C. Ernhart, M. Craft, P. Berman. Learning of Relative And Absolute Size Concepts in Preschool Children. – *Journal of Experimental Child Psychology*, Vol. **1**, No 1, 1964, 26-36.
7. Hagmann-von Arx, P., C. Meyer, A. Grob. Assessing Intellectual Giftedness With The WISC-IV And The IDS. – *Zeitschrift für Psychologie*, Vol. **216**, No 3, 2008, 172-179.
8. Kaufman, A.. *Intelligent Testing With the WISC-III*. John Wiley & Sons, 1994.
9. Kemp, C., J. Tenenbaum. The Discovery of Structural Form. – *Proceedings of the National Academy of Sciences*, Vol. **105**, No 31, 2008, 10687-10692.
10. Lillard, A.. *Montessori: The Science Behind the Genius*. Oxford University Press, USA, 2008.
11. Lillard, A., N. Else-Quest. The Early Years: Evaluating Montessori. – *Science*, Vol. **313**, No 5795, 2006, 1893-1894.
12. Montessori, M.. *The Montessori Method*. Frederick A. Stokes Co., 1912.
13. Nakamura, T., T. Nagai, N. Iwahashi. Multimodal Object Categorization by a Robot. – In: *Proceedings of the IEEE/RSJ IROS*, 2415-2420.
14. Natale, L., G. Metta, G. Sandini. Learning Haptic Representation of Objects. – In: *Proc. of the Intl. Conf. on Intelligent Manipulation and Grasping*.
15. Nolfi, S., D. Marocco. Active Perception: A Sensorimotor Account of Object Categorization. – In: *From Animals to Animats: 7*, 266-271.
16. Pitamic, M.. *Teach Me to Do It Myself*. Elwin Street Productions, 2004.
17. Sinapov, J., T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, A. Stoytchev. Interactive Object Recognition Using Proprioceptive And Auditory Feedback. – *The Intl. J. of Robotics Research*, Vol. **30**, No 10, 2011, 1250-1262.
18. Sinapov, J., A. Stoytchev. Detecting the Functional Similarities Between Tools Using a Hierarchical Representation of Outcomes. – In: *Proceedings of the 7th IEEE International Conference on Development and Learning (ICDL)*, 91-96.
19. Sinapov, J., A. Stoytchev. The Odd One Out Task: Toward an Intelligence Test for Robots. – In: *Proceedings of the 9th IEEE International Conference on Development and Learning (ICDL)*, 126-131.
20. Sinapov, J., A. Stoytchev. Object Category Recognition by a Humanoid Robot Using Behavior-Grounded Relational Learning. – In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 184-190.
21. Sinapov, J., V. Sukhoy, R. Sahai, A. Stoytchev. Vibrotactile Recognition and Categorization of Surfaces by a Humanoid Robot. – *Robotics, IEEE Transactions on*, Vol. **27**, No 3, 2011, 488-497.
22. Sugarman, S.. The Cognitive Basis of Classification in Very Young Children: An Analysis of Object-Ordering Trends. – *Child Development*, Vol. **52**, No 4, 1981, 1172-1178.
23. Sun, J., J. Moore, A. Bobick, J. Rehg. Learning Visual Object Categories for Robot Affordance Prediction. – *The Intl. J. of Robotics Research*, Vol. **29**, No 2-3, 2010, 174.
24. Takamuku, S., K. Hosoda, M. Asada. Object Category Acquisition by Dynamic Touch. – *Advanced Robotics*, Vol. **22**, No 10, 2008, 1143-1154.
25. Tenenbaum, J., V. De Silva, J. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. – *Science*, Vol. **290**, No 5500, 2000, 2319-2323.