Computational Benefits of Social Learning Mechanisms: Stimulus Enhancement and Emulation

Maya Cakmak, Nick DePalma, Rosa Arriaga, Andrea L. Thomaz
Center for Robotics and Intelligent Machines
Georgia Institute of Technology
801 Atlantic Dr., Atlanta GA 30332
{maya, ndepalma, arriaga, athomaz}@cc.gatech.edu

***Abstract*—Social learning in robotics has largely focused on imitation learning. In this work, we take a broader view of social learning and are interested in the multifaceted ways that a social partner can influence the learning process. We implement *stimulus enhancement* and *emulation* on a robot, and illustrate the computational benefits of social learning over individual learning. Additionally we characterize the differences between these two social learning strategies, showing that the preferred strategy is dependent on the current behavior of the social partner. We demonstrate these learning results both in simulation and with physical robot 'playmates'.**

I. Introduction

Social partners can guide the learning process by directing a learner's attention to informative parts of the environment or by suggesting informative actions for the learner. Humans and some animals are equipped with several mechanisms that take advantage of social partners. Understanding these mechanisms and their role in learning will be useful in building robots with similar abilities. Such robots can maximally benefit from other agents (humans or robots) in their environment, as well as explicit teaching attempts by these agents.

Our research is motivated by the four social learning mechanisms identified in biological systems [Tomasello, 2001], [Call and Carpenter, 2002]:

- *Stimulus (local) enhancement* is a mechanism through which the attention of an observer (youngster or novice) is drawn to objects that others are interacting with. This facilitates learning by focusing the exploration of the observer on interesting objects – objects that have useful affordances for other members of the social group.
- *Emulation* is a process in which the observer witnesses a particular result on an object while others interact with it, but then employs its own action repertoire to produce the same result on the same object. In this case, learning is facilitated both by attention being directed to the object of interest and by the observation of the goal.
- *Mimicking* corresponds to the observer copying the actions of others without an appreciation of their purpose (goal or intention). The observer later comes to discover the effects of the action in different situations. Mimicking helps the observer by suggesting actions that can produce useful results.
- *Imitation* refers to reproducing the actions of others to obtain the same results with the same goal.

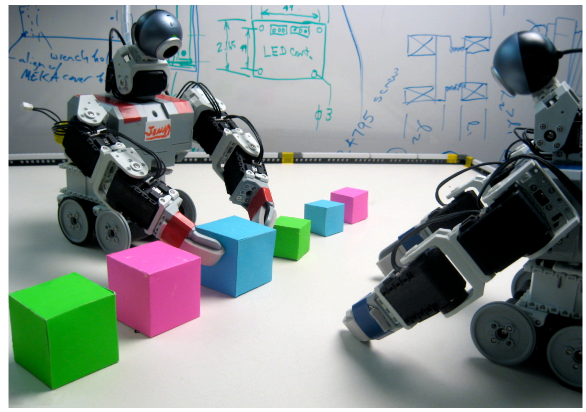

Fig. 1. Robot playmates Jimmy and Jenny in the playground.

Social learning in robotics has often focused on the last and most complex of these four mechanisms–imitation. Most strive to create robots capable of imitative behavior [Shaal, 2002], resulting in the ability to reproduce demonstrated motor actions. Through imitative behavior the robot can learn generalized task representations [Pardowitz and Dillmann, 2007], policies [Chernova and Veloso, 2007] or a proto-language about actions [Billard, 2002]. This is called *Learning by Imitation*. A body of research has been devoted to finding ways to learn the imitative behavior itself, rather than hard-coding it [Nehaniv and Dautenhahn, 2002], [Demiris and Hayes, 2002]. Others focus on environmental aspects of a demonstration, making the robot reproduce the task without imitative behavior, but instead with it's own actions [Kuniyoshi et al., 1994], [Montesano et al., 2008], [Jansen, 2005]. This is known as task- or goal-level imitation and resembles emulation. In a closely related field, *Learning by Demonstration*, researchers focus on adjusting a robot's *actions* while imitating a human demonstration. This can involve extracting trajectory representations in joint and task space [Calinon and Billard, 2008], dynamical equations to control the robot's movements [Pastor et al., 2009] or a sequence of primitive actions [Amit and Mataric, 2002].

Alternatively, we take a broader view of social learning and are interested in the multifaceted ways that a social partner can influence the learning process. In this study we implement *stimulus enhancement* and *emulation* on a robot. The goal of our research is to assess the computational benefits of social learning over individual exploration. Exploration is the process of collecting experiences of interactions with the environment. In terms of robot learning, these social learning mechanisms can be viewed as ways of guiding the robot's exploration of the learning space. Both stimulus enhancement and emulation direct the attention of the learner to more informative or salient parts of the *feature space* (*i.e.* the environment).

We show that both stimulus enhancement and emulation provide learning benefits over individual exploration, particularly in the case when the target goal of learning is a rare occurrence in the environment. We analyze the effects of using different strategies in environments with different affordance rareness and different social partners. We also characterize the difference between the two social strategies, and show that each is better than the other depending on the current behavior of the social partner. We demonstrate these learning results in both simulation and with two physical robot 'playmates'.

II. Approach

In this work, we have a social learning situation composed of two robot playmates with similar action and perception capabilities. Our experiments are focused on a robot's performance in the task of learning the sound-making affordance of different objects present in the environment.

A. Robot Platform

We use two robots, Jimmy and Jenny (Fig. 1), which are upper torso humanoids with wheels, built from Bioloid kits and a Webcam. They are approximately 10 inches high and have 8 degrees of freedom, which enables arm movements, torso rotation and neck tilt. The wheels are used for navigating the workspace.

The playmates' behavior system is implemented in C6, the latest revision of the *Creatures* cognitive architecture for interactive characters [Blumberg et al., 2002]. The same behavior system is used to control the real robots with percepts obtained from the real sensors, as well as a graphical model of the robots with simulated sensing and world dynamics. The simulation allows setting up different environment compositions with different object properties.

The behavior system implements a finite state machine to control the exploration of the robot for collecting interaction experiences. In individual exploration the robot (i) observes the environment, (ii) approaches the most salient object, (iii) performs the selected action, (iv) observes the outcome of the interaction (sound or no sound), (v) goes back to its initial position and (vi) updates the saliency of objects and the desirability of actions based on its exploration strategy. In social exploration, after every object interaction the robot goes into an *observing* state and performs the same updates, of object saliency and action desirability, based on its observation of the other agent's interaction in the environment.

The robots communicate the object that they interact with and the parameters of the action they used in the interaction to the other robot through network messages. The observation of the other robot's interaction is based on these messages and the sound perceived during the observation state. The sound perception is based on volume thresholding on the signal obtained through the microphone embedded on the Webcams.

B. Learning Task

Our experiments focus on the task of learning object affordances. Affordance learning is formulated as learning a relation between a *context* in which an *action* is performed to produce a certain *outcome*. The relation is learned from interaction experiences consisting of [context; action; outcome] tuples [Sahin et al., 2007]. We use a 2-class Support Vector Machine (SVM) classifier[1] that predicts the effect of an action in a given environmental context. The input to the SVM consists of the perceived features of the interacted object and the parameters of the action performed on the object. The prediction target is whether or not this interaction produces sound. In this framework the robot is simultaneously learning the object features and action parameters required to produce a desired effect in the environment.

Our goal is to compare social and individual *exploration strategies*, *i.e.* the set of rules adopted for interacting with the environment to gain experience for learning affordances. An exploration strategy is implemented as an attention mechanism, where each object attribute or action parameter has a corresponding saliency and the robot interacts with the most salient object by performing the most salient action. Each strategy has a different rule for updating saliencies after every interaction. While individual exploration can take into account past experiences, social exploration can also benefit from the observed interactions of the other robot.

C. Objects

The learning environment involves objects that are perceived through three discrete attributes: *color*, *size* and *shape*. Each object also has the hidden property of *sound-makingness*. Different environments are set up by choosing objects with different combinations of attributes, and different sound-making properties. For instance, all green objects could be sound makers in one environment, while in another all objects with a particular shape and size larger than a threshold are sound-makers.

In our experiments we systematically vary the frequency of sound-makers in the environment and compare various individual and social exploration strategies. We hypothesize, based on prior work, that social learning will be especially beneficial in the case of rare sound-makers [Thomaz and Cakmak, 2009].

The simulation environment includes 24 objects with different attributes (one of 4 colors, 3 sizes and 2 shapes). We control the percentage of objects in the environment that are able to produce sound, resulting in six different learning

[1]The choice of classifier is not crucial for the results of this study. SVMs are widely used discriminative classifiers.

environments with 75%, 50%, 25%, 17%, 8%, and 4% sound-makers. The physical experiments have a simplified environment with 8 objects (4 colors and 2 sizes). These experiments are performed in two learning environments where (i) all small objects are sound-makers (50%) and (ii) only one object is a sound-maker (12%).

D. Actions

The playmates' action set consists of two actions: *poke*–a single arm swing (*e.g.*, for pushing or rolling objects) and *grasp*–a coordinated swing of both arms. Both actions involve an initial phase of *approaching* an object of interest. Both actions are parametrized such that different versions of the actions are obtained with different (i) acting distances, (ii) grasp widths and (iii) poking speeds. The actions are tuned to have different effect on each object to set up different learning problems. All action parameters are discrete and there are 24 different actions (poke or grasp, 4 grasp width values, 4 poke speed values and 3 acting distances). On the physical robots there are 8 possible actions (2 possible values for each action parameter).

As we do with objects, we can vary the frequency of sound-producing interactions by controlling the percentage of actions that produce sound. We set the environment and the actions in the simulation experiments such that 25% of all possible actions is able to produce sound when executed on a sound-maker. This is achieved by making only the *grasp* action produce sound when the grasp width is smaller than a threshold and the grasp is performed at a certain distance from the objects. In the physical experiment this results in 12% of actions being a sound producing one.

III. Experiments

We have two sets of experiments to address individual vs. social learning. Since our goal is to compare the social learning strategies to individual learning, we first collect performance data for individual learning. However, a question arises as to what is a fair or appropriate individual learning baseline to compare against. We consider three different individual exploration strategies for learning affordances: *random*, *goal-directed* and *novelty-based* exploration. We also compare these strategies with a *systematic* data set that consists of all possible interactions.

Second, we present experiments evaluating social learning via stimulus enhancement and emulation. In this experiment one of the robots (*learner*) explores the environment using these social strategies while the other robot (*social partner*) interacts with the environment with a pre-defined preference. The preference determines how much the learner can benefit from the social partner as a 'teacher'.

A. Individual Learning

Three individual exploration strategies are implemented.

1) Random Exploration: In each interaction the robot randomly picks a new object, action and a set of action parameters. This is achieved by randomizing the saliency of each object attribute and action parameter and selecting the most salient object and action. The data sets collected with random exploration are equivalent to random subsets of the systematic data set.

2) Goal-directed Exploration: In goal-directed exploration, the robot keeps interacting with objects similar to ones that have given the desired effect in a previous interaction. Similarly it performs actions that are similar to those that produced sound in the past. If an interaction produces sound, the saliency of some attributes of the object used in that interaction are increased and the saliency of different ones are decreased. Increasing or decreasing all attributes deterministically is avoided because this will result in interacting with the exact same object once it has produced sound, therefore will stop the exploration. By updating a random subset of the attributes of an object that made sound, the robot will interact with objects that have common attributes, rather than exactly the same object. If no sound is produced the robot keeps updating saliencies randomly.

3) Novelty-based Exploration: In this strategy the robot prefers novel objects and actions. After every interaction the robot reduces the saliency of attributes of the object that it interacted with, while increasing the saliency of different attributes. Actions and action parameters are altered similarly.

4) Systematic Data Set: In addition to the three exploration strategies, we consider a data set that consists of the complete the learning space. This is collected by going over all possible object-action combinations, one-by-one. In the simulation experiment this results in a data set with 576 interaction tuples and in the physical experiment it requires 64 interactions. The three exploration strategies explained above are compared on smaller data sets.

B. Social Learning

The second set of experiments determine the benefits of social exploration strategies. As in the individual experiments, we systematically vary the frequency of sound-makers in the environment. The number of interactions is fixed to 58 for all strategies.

The social partner has a crucial role in social learning. In order to assess the effect of the social partner's behavior on the learner we vary its behavior systematically. The social partner can give three types of demonstrations:

- *Goal-demonstration*: The learner's target goal (sound) is demonstrated by interacting with a sound-maker object using a sound-producing action.
- *Object-demonstration*: The demonstration involves a sound-maker object but sound is not actually produced, *i.e.* a non-sound-producing action is used.
- *Negative-demonstration*: The demonstration involves a non-sound-maker object.

TABLE I
DEMONSTRATION TYPE PREFERENCES FOR THREE SOCIAL PARTNER BEHAVIORS.

Demo. Type	Same-goal	Same-obj.	Different-goal
Goal-demo.	60%	20%	20%
Object-demo.	20%	60%	20%
Neg.-demo.	20%	20%	60%

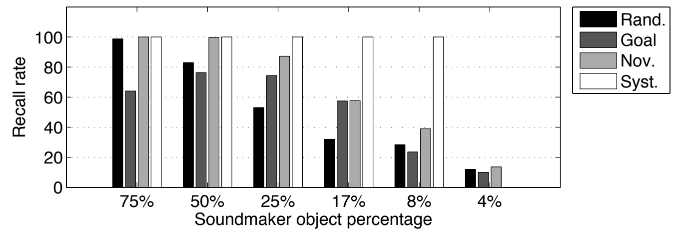

Fig. 2. Recall rate for individual learning strategies after 58 interactions for six environments with different sound-maker frequencies.

We investigate three social partner behaviors each of which provides a particular type of demonstration more than others. The percentage of each type of demonstration in these behaviors is given in Table I.

- Social partner with *same goal* frequently demonstrates the learning target of the learner (sound-making).
- Social partner with *same object interests* spends most time performing object-demonstrations.
- Social partner with *different goal* interacts mostly with objects that are not sound-makers, therefore providing many negative-demonstrations.

We implement two social exploration strategies. Both strategies influence the way that object attribute saliencies are updated but do not influence action selection. In order to avoid superimposed effects, action selection is random in social exploration strategies.

1) Stimulus Enhancement: In this strategy, the robot prefers to interact with objects that its playmate has interacted with. After every observed interaction, the learner increases the saliency of attributes of the object that the social partner has interacted with and decreases the saliency of other attributes.

2) Emulation: In this strategy, the robot prefers to interact with objects that have given the desired effect during its social partner's interactions. If an observed interaction produces sound, the saliencies of the attributes of the interacted object are increased. Otherwise, the saliencies are randomized.

IV. RESULTS

We first present results from simulation for all environments, exploration strategies and social partner behaviors. Then in a simplified environment we present a comparison of individual and social learning strategies on the physical robots. Our primary performance measure is recall rate[2] in prediction using the complete set of systematic interactions as a test set.

A. Baseline: Results for Individual Learning

1) Comparison of strategies: Systematic exploration is designed to cover the complete learning space. All other strategies are tested on the systematic data set. The systematic strategy is a best-case scenario for the learning algorithm, and essentially shows that this is a learnable problem.

A 20-fold cross validation test is performed on the systematic data set for the 6 environments described in Section II-C. We observe that prediction is 100% accurate for the systematic strategy in all environments with sound-maker frequency of 8% or greater. In the last environment case (4% sound-makers) the event of sound-making happens so infrequently that the resulting SVM always predicts 'no sound' and the recall rate is 0%.

Fig. 2 compares the recall rate for individual learning strategies in different environments. Random, goal-directed and novelty-based exploration strategies are repeated 50 times in each environment with randomized initialization. Their performance is compared for learning with 58 interaction samples, which corresponds to 10% of the systematic data set.

The performance of the random exploration strategy reduces as the sound-maker objects become rare in the environment, since it is less likely to randomly interact with a sound-maker when it is rare.

The goal-directed strategy results in lower recall rates than random when the sound-makers are frequent in the environment. With this strategy the robot interacts only with a subset of objects that are similar to the first object that was discovered to be a sound-maker. However, when the environment has a high percentage of sound-makers, objects with no common perceptual attributes can all be sound-makers. Therefore, in such environments covering only a subset of objects degrades the performance of the goal-directed strategy. As the sound-makers become less frequent the goal-directed strategy become better than random the random strategy.

The novelty-based strategy outperforms the other exploration strategies especially when the sound-makers are frequent. The strength of this strategy in these environments is its uniform coverage of the search space by always interacting with different objects. As the sound-makers become very rare the performance of all three strategies degrade and the difference between the strategies becomes less significant.

2) Comparison of environments: In Fig. 2, we see a significant effect of the rareness of sound-makers in the environment on all three exploration strategies (see Table II for statistical significance). While the performance of random and novelty-based strategies monotonically decrease with decreasing sound-maker frequency, the performance of the goal directed strategy increases initially and decreases afterwards for reasons explained previously.

[2]Recall corresponds to the ratio of true positives and the sum of true positives and false negatives. Due to space limitations, in this paper we restrict our analysis to effects on recall rate.

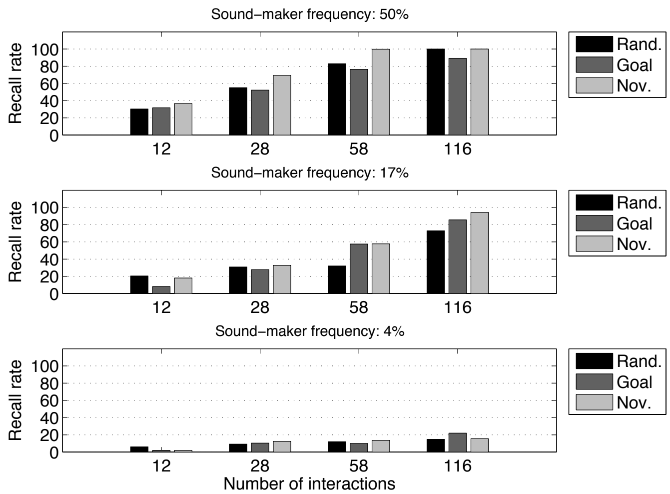

Fig. 3. Recall rate for individual learning strategies after different number of interactions for three sample environments with different sound-maker frequencies.

TABLE II

EFFECT OF SOUND-MAKER RARENESS ON DIFFERENT EXPLORATION STRATEGIES (MEASURED WITH 1-WAY ANOVA).

Strategy	Analysis of variance
Random	$F(5, 294) = 57.54$, $p < .001$
Goal-directed	$F(5, 294) = 27.48$, $p < .001$
Novelty-based	$F(5, 294) = 79.87$, $p < .001$
Stimulus enh.	$F(5, 294) = 1.10$, $p > .05$
Emulation	$F(5, 294) = 2.08$, $p > .05$

3) Comparison of number of interactions: All three strategies result in imperfect learning because they cannot explore the complete space of possible objects and actions. However, we expect that the longer we allow the robot to interact with the environment, the better its learning will be. In Fig. 3 we compare the performance of learning with 12, 28, 58 and 116 samples collected with each exploration strategy. These respectively correspond to 2%, 5% 10% and 20% of all systematic interactions. In general, it can be observed that learning improves with increasing number of interactions. However, in the case of very rare sound-makers (4%) increasing the number of interactions does not improve performance. The reason is that when the sound-makers are very rare, even more than 116 interactions are necessary to randomly discover a sound-maker.

B. Social versus Individual Learning

Next we compare prediction performance for learning with social versus individual exploration strategies in the same environments, considering the *same goal* case of the social partner (Fig. 4). With each strategy the robot learns from 58 interactions and 50 trials are performed with random initialization for each case.

First, we observe that in all environments social learning strategies result in better learning and the performance gain is more pronounced in rare affordance cases. The effect of

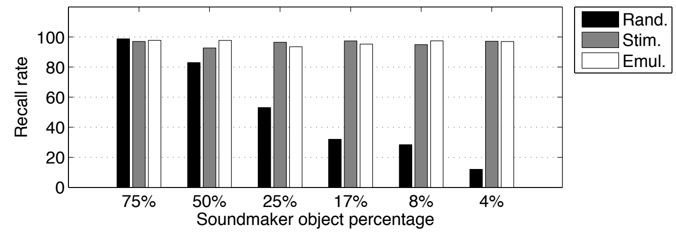

Fig. 4. Comparison of recall rate in individual (random) and social learning for 6 different environments with different sound-maker rareness. The social partner behavior is *same goal*.

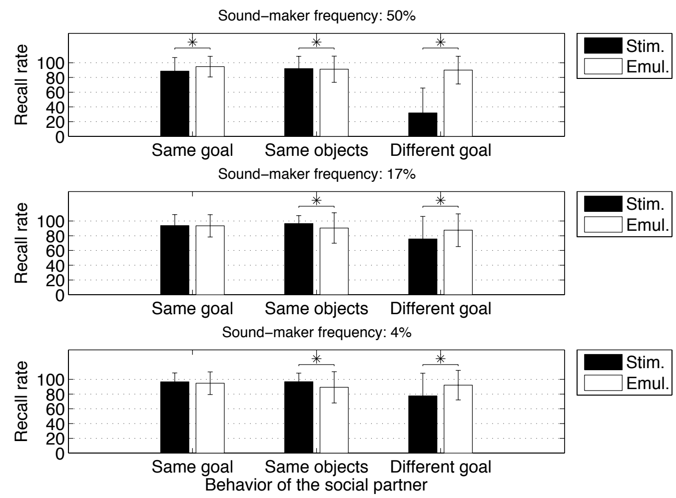

Fig. 5. Comparison of stimulus enhancement and emulation for different social partner behaviors ('*' indicates a significant difference: i.e., a T-test comparison between the two groups of data results in $p < .05$).

rareness on stimulus enhancement and emulation based learning is very small and is not statistically significant (Table II). It can also be observed in Fig. 4 that the performance of social learning strategies stay high as rareness increases. This supports our hypothesis that social learning is particularly beneficial in learning rare affordances.

C. Stimulus Enhancement versus Emulation

We compare stimulus enhancement and emulation in three cases where the behavior of the social partner is altered as described in Section III-B. Fig. 5 gives a comparison in three environments for learning with 58 interactions and 200 random initializations.

In the case of a same goal partner, most of the interactions with sound-maker objects produce sound. This results in both strategies modifying object attribute saliencies in similar ways. Thus, they have similar performances in this case.

In the case of a partner with the same object interests, we observe that stimulus enhancement results in better learning performance. In this case most of the interactions of the social partner are with a sound-maker, causing the stimulus enhancement strategy to focus its exploration on the right objects. On the other hand these interactions are often not informative for the emulation strategy since sound is not produced. Consequently, the emulation strategy reduces to random exploration, resulting in a decrease in performance.

Finally, in the case of a partner with different goals, we observe that emulation outperforms stimulus enhancement. In this case both strategies cannot benefit from the social partner 60% of time since it is interacting with non sound-makers. However, the stimulus enhancement strategy is biased to interact with those objects and cannot discover sound-makers. The emulation strategy on the other hand randomly explores the environment when the partner's action produces no sound and has more chance to discover the sound-makers.

The differences in performance of stimulus enhancement and emulation is always significant in the two cases where the social partner does not share the goal of the learner. The performance of stimulus enhancement is especially affected in the environment with high number of sound-makers, since the few interactions in which the social partner demonstrates the sound-makers are not sufficient to cover a large portion of the sound-makers.

D. *Validation on the Physical Robots*

A simplified version of the simulation experiments was run on the physical robots as described in Section III. Table III gives the results of learning for two different environments and three strategies. The results support our findings from the simulation experiment: (i) social exploration strategies yield better learning of affordances and (ii) increasing rareness has less effect on social learning strategies.

Additionally, Fig. 6 gives the progress of attribute saliencies during the first 16 interactions in emulation based exploration. All small-sized objects make sound in this environment (50% sound-maker). It can be observed that when the social partner's action emits sound, the saliencies of the attributes of the sound-producing object are increased, while others are decreased. For instance between the 6th and 7th interactions of the learner, no sound is produced by the demonstrator and *large* size becomes more salient than *small* due to random exploration. Between the 7th and 8th interactions, a sound is caused by the demonstrator and small size becomes most salient for the 8th interaction. As a result of the emulation based exploration the saliency of the size attribute is higher for *small* in the majority of the experiment, and the learner has more chance to discover the actions that produce sound by interacting with small objects. On the other hand, we do not observe any trends in the saliencies of the color attribute, and the interactions are equally distributed over different colors.

V. Discussion

As expected from prior work, social learning is better than individual learning especially in cases where the learned affordance is rare. In this work we've shown the computational

TABLE III
RECALL RATE IN PHYSICAL ROBOT EXPERIMENTS.

	50% Sound-maker	12% Sound-maker
Rand.	75%	0%
Stim.	100%	100%
Emul.	100%	100%

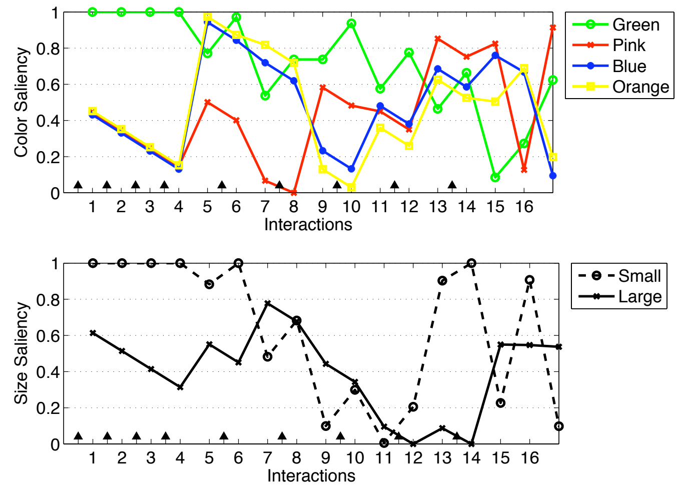

Fig. 6. Progress of saliencies over interactions in emulation based exploration in the robot experiment in which all small objects are sound-makers. The interactions before which the demonstrator emits sound are shown with arrows. We see that emulation causes small size to have increased saliency, whereas it the color features don't show similar trends.

benefits of two biologically inspired social learning mechanisms: stimulus enhancement and emulation.

We find that both have a positive impact on learning, particularly when the learning goal has a relatively rare occurrence in the given environment. More intelligent individual strategies, such as novelty-based exloration, improve the data set by taking into account previous interactions. However the discovery of interesting regions of the search space is still impacted by rareness. This weakness of individual learning is exactly where social learning can help. Social partners can demonstrate rare affordances or point the learner to relevant objects such that useful information can be acquired right away, without spending time discovering the right objects.

Stimulus enhancement is better than emulation in situations where the social partner interacts with useful objects without necessarily demonstrating the affordances. This may occur in cases where an object has different affordances for the observer and the social partner due to different motor/perceptual capabilities. This is often the case in parent-child interactions.

Emulation is better than stimulus enhancement in situations where social partners have different interests. This can be a situation in which the social partner has goals that do not involve the affordance in which the learner is interested in, or when the learner does not have the perceptual capabilities to observe the effects of the its partner's actions (and therefore cannot adopt those effects as its goal). In these cases, stimulus enhancement results in paying attention to objects that do not have the affordance of interest, while emulation can capitalize on the few demonstrations given by social partners and keep searching on their own in the rest of the time.

In this study action selection is random during social exploration. We are currently extending social learning mechanisms to action selection as well. This involves an exploration of the action space guided by the social partners. For instance the robot can copy its social partner's actions in every interaction or when a sound is produced by the action. It can also combine object exploration strategies with these action exploration strategies. Such mechanisms would correspond to *mimicking* and *imitation* as described in Section I.

VI. Conclusion

We presented a series of experiments on two social learning mechanisms: stimulus enhancement and emulation. We looked at the task of learning the sound-making affordance of different objects by a robot, while another robot (a social partner) is also interacting with the same objects. Our experiments support that both strategies provide learning benefits over individual exploration. We characterized the difference between these two strategies, showing that there are certain situations where one can be better than the other. We also investigated the effects of affordance rareness and number of interactions on learning.

What we can draw from this work is that each social learning strategy has its own purpose and is beneficial in different ways, which is not surprising from a developmental perspective. The contribution of this work is the articulation of the computational benefit of these social learning strategies for a robot learner. The fact that each strategy performs differently in different situations indicates the importance for the robot to have all of these strategies available. This would allow the robot to switch learning strategies in case one particular strategy is not leading to useful learning experiences in a particular environment. Other environmental (or social) cues may also indicate to the learner which strategy is the right one to use in a particular context.

References

[Amit and Mataric, 2002] Amit, R. and Mataric, M. (2002). Learning movement sequences from demonstration. In *Proceedings of the IEEE 1st Intl. Conference on Development and Learning*, pages 203–208.

[Billard, 2002] Billard, A. (2002). Imitation: a means to enhance learning of a synthetic protolanguage in autonomous robots. In Dautenhahn, K. and Saunders, J., editors, *Imitation in animals and artifacts*. MIT Press.

[Blumberg et al., 2002] Blumberg, B., Downie, M., Ivanov, Y., Berlin, M., Johnson, M., and Tomlinson, B. (2002). Integrated learning for interactive synthetic characters. In *Proc. of the ACM SIGGRAPH*.

[Calinon and Billard, 2008] Calinon, S. and Billard, A. (2008). A probabilistic programming by demonstration framework handling skill constraints in joint space and task space. In *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*.

[Call and Carpenter, 2002] Call, J. and Carpenter, M. (2002). Three sources of information in social learning. In Dautenhahn, K. and Saunders, J., editors, *Imitation in animals and artifacts*. MIT Press.

[Chernova and Veloso, 2007] Chernova, S. and Veloso, M. (2007). Multiagent collaborative task learning through imitation. In *Proceedings of the fourth International Symposium on Imitation in Animals and Artifacts*.

[Demiris and Hayes, 2002] Demiris, J. and Hayes, G. (2002). Imitation as a dual-route process featuring predictive and learning components: a biologically-plausible computational model. In Dautenhahn, K. and Saunders, J., editors, *Imitation in animals and artifacts*. MIT Press.

[Jansen, 2005] Jansen, B. (2005). Goal level imitation with emergent repertoires. In *Proceedings of the third Intl. Symposium on Imitation in Animals and Artifacts*, pages 62–68, University of Hertfordshire, Hatfield, England.

[Kuniyoshi et al., 1994] Kuniyoshi, Y., Inaba, M., and Inoue, H. (1994). Learning by watching: Extracting reusable task knowledge from visual observation of human performance. *IEEE Transactions on Robotics and Automation*, 10:799–822.

[Montesano et al., 2008] Montesano, L., Lopes, M., Bernardino, A., and Santos-Victor, J. (2008). Learning object affordances: From sensory-motor coordination to imitation. *IEEE Transactions on Robotics*, 24:15–26.

[Nehaniv and Dautenhahn, 2002] Nehaniv, C. and Dautenhahn, K. (2002). The correspondence problem. In Dautenhahn, K. and Saunders, J., editors, *Imitation in animals and artifacts*. MIT Press.

[Pardowitz and Dillmann, 2007] Pardowitz, M. and Dillmann, R. (2007). Towards life-long learning in household robots: The piagetian approach. In *Proceedings of the IEEE 6th Intl. Conference on Development and Learning*.

[Pastor et al., 2009] Pastor, P., Hoffmann, H., Asfour, T., and Schaal, S. (2009). Learning and generalization of motor skills by learning from demonstration. In *IEEE Intl. Conference on Robotics and Automation*.

[Sahin et al., 2007] Sahin, E., Cakmak, M., Dogar, M., Ugur, E., and Ucoluk, G. (2007). To afford or not to afford: A new formalization of affordances toward affordance-based robot control. *Adaptive Behavior*, 15(4):447–472.

[Shaal, 2002] Shaal, S. (2002). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(11):233–242.

[Thomaz and Cakmak, 2009] Thomaz, A. and Cakmak, M. (2009). Learning about objects with human teachers. In *HRI '09: Proceedings of the ACM/IEEE Intl. Conference on Human-Robot Interaction*.

[Tomasello, 2001] Tomasello, M. (2001). *The Cultural Origins of Human Cognition*. Harvard University Press.