

Learning by Imitation, Reinforcement and Verbal Rules in Problem Solving Tasks

Frédéric Dandurand, Melissa Bowen, Thomas R. Shultz

McGill University, Department of Psychology, 1205 Dr. Penfield Ave., Montréal, Québec, H3A 1B1, Canada, fdandu@ego.psych.mcgill.ca, melissa.bowen@mail.mcgill.ca, thomas.shultz@mcgill.ca

Abstract

Learning by imitation is a powerful process for acquiring new knowledge, but there has been little research exploring imitation's potential in the problem solving domain. Classical problem solving techniques tend to center around reinforcement learning, which requires significant trial-and-error learning to reach successful goals and problem solutions. Heuristics, hints, and reasoning by analogy have been favored as improvements over reinforcement learning, whereas imitation learning has been regarded as rote memorizing. However, research on imitation learning in animals and infants suggests that what is being learned is the overall arrangement of actions (sequencing and planning). Applied to problem solving, this suggests that imitation learning might enable a problem solver to infer a complex hierarchical problem representation from observation alone.

We compared three types of learning in problem solving tasks: imitation learning (a group that viewed successful problem solving demonstrations), reinforcement learning (a group that got feedback indicating whether their answer was correct or not) and explicit rule learning (a group that was presented specific instructions to solve the problem). The task required participants to find, with three uses of a scale, the one ball which was either heavier or lighter than the rest of a set of 12 balls. We found that subjects in the imitation learning and explicit learning groups outperformed those in the reinforcement learning group. We conclude that learning by imitation in problem solving tasks is worthwhile, efficient and even superior to explicit learning because of the minimal time and energy investment required from the mentor.

1. Introduction

Problem solving is “thinking that is directed toward the solving of a specific problem that involves both the formation of responses and the selection among possible responses” [25]. It is therefore a very important area of cognitive psychology, and it is considered a crucial com-

ponent of intelligence. “The ability to solve problems is one of the most important manifestations of human thinking.” [12].

Information-processing theory is currently the dominant approach to problem solving [12]. Problems are construed in terms of states, transitions and operators. The essence of problem solving is a search through the state space for a solution state using the operators available at each point while satisfying a set of problem-specific constraints.

1.1. Learning in Problem Solving Tasks

Several types of learning can occur during problem solving efforts, either in isolation or in various combinations.

1.1.1. Reinforcement Learning

Reinforcement learning, or learning by trial-and-error, is a mechanism whereby systems (humans, animals or machines) seek to maximize rewards and minimize punishments. In general, rewarded behaviors tend to increase in frequency whereas punished behaviors tend to decrease in frequency as learning progresses [2, 16]. Applied to problem solving, reinforcement learning allows problem solvers to learn which tactics are successful and which ones are not using feedback: rewarding searches that yield correct solutions and punishing those that do not. It is the most common learning mechanism in problem solving because feedback information is usually embedded in the problem itself or is available in the environment. The main problem with learning by reinforcement is that the information available through feedback is typically very limited, in the form of binary data (correct/satisfactory or incorrect/unsatisfactory answer).

Although he used a different terminology, Holyoak describes a form of reinforcement learning as the learning mechanism involved in information-processing theory: “An intelligent problem solver uses the results of solution attempts to acquire new knowledge that will help solve similar problems more readily in the future” [12].

1.1.2. Explicit Learning

Explicit learning, or verbal learning, is based on having access to explicit instructions for solving problems. They are typically expressed as abstract symbolic rules in the form of *if... then...* statements. Western cultures particularly value explicit verbal instructions such as problem solving algorithms typically described in textbooks, documents and other “how-to” manuals.

Explicit learning is limited in scope. First, it assumes the availability of a skilled teacher who has the time, energy and ability to express problem solving reasoning explicitly, concisely, completely and coherently.

Second, broad classes of problems, such as information-integration category-learning tasks [3], are highly problematic for explicit instructions. Even skilled problem solvers are unable to express their problem solving strategies explicitly because solution strategies for these problems are learned and accessed implicitly.

1.1.3. Imitation Learning

Imitation learning, rooted in the long tradition of social learning [4], can be defined as a mechanism where behaviors or skills are acquired by watching others perform. “Perhaps the most important learning technique in the social domain is that of imitation, or observational learning” [24]. In the natural world, learning by imitation makes evolutionary sense for social animals because it allows them to learn and transmit successful methods and strategies, possibly acquired over many generations.

To date, there is still a lot of controversy around a precise and detailed definition of imitation learning. Definitions can generally be regarded on a spectrum from inclusive to restrictive. For instance, Thorndike (1898) defined imitation as any situation in which animals “from an act witnessed learn to do an act.” [1] In contrast, Thorpe (1963) defined “true imitation as the copying of a novel or otherwise improbable act or utterance, or some act for which there is clearly no instinctive tendency” [1].

Most of the controversy originates from the fact that other mechanisms (generally based on some kind of priming) can also account for imitative behaviors. Those mechanisms include [1, 17]:

1. Social Facilitation / Social Enhancement - The mere presence of conspecifics encourages similar behaviors.
2. Local Enhancement - The attention of the observer is drawn to a place or location due to activities of the demonstrator.
3. Stimulus Enhancement - The attention of the observer is drawn to an object (e.g., tool) due to activities of the demonstrator.
4. Goal Emulation - The imitator does not try to copy the action, but tries to reproduce the result.

Most imitation researchers agree about the importance of ruling out these simpler mechanisms. To truly qualify as imitation, it appears that some kind of understanding of the demonstrator’s intentions is important. Recent research has found evidence of such understanding of intentions in animals and in human infants [8, 9, 27].

Byrne and Russon [7] proposed the idea of a *program level* imitation consisting in imitating the overall arrangement of actions in a hierarchical fashion, particularly the planning of and sequencing of actions. They contrasted this with *action level* imitation where the fine details of the actions are copied or imitated. They argued that program level imitation qualifies as “true imitation” because it implies that the imitator understands the intentions of the demonstrator in terms of goals and sub-goals. They consider the learning of a new arrangement of behavioral units already present in the behavioral repertoire qualifies as novel.

In the context of problem solving, imitation learning takes the form of demonstrations. Byrne and Russon’s theory suggests that imitation learning might enable a problem solver to infer a complex hierarchical problem representation from observation alone.

Learning by imitation does not have the limitations of explicit learning. Because it conveys information or instructions implicitly in problem solving demonstrations, it can be applied to tasks learned implicitly and to tasks for which no written instructions can be found. In fact, the demonstrator need not be conscious he is being imitated, although awareness of his role might facilitate learning because he can emphasize or highlight critical steps during the demonstration. This makes learning by imitation more efficient and adaptive than explicit learning in many contexts.

Learning by imitation and by reinforcement are probably mediated by different brain mechanisms, namely mirror neurons for imitation learning [18] and the activity of mesencephalic dopaminergic neurons for reinforcement learning [14]. Furthermore, both types of learning are used in machine learning [6, 26].

1.2. Supplemental Approaches to Problem Solving

Several methods for either complementing/supplementing reinforcement learning or for replacing it have been proposed. Techniques traditionally favored include: teaching of heuristics [20], hints [15] and reasoning by analogy [13]. Holyoak [12] proposes that heuristics are used to limit search complexity by considering only a small number of alternatives that seem most likely to lead to a solution.

In contrast, the importance of learning by demonstration (imitation) has been minimized because it was considered as rote memorizing [15], and therefore trivial and

uninteresting. In the light of the more recent research just reviewed, we believe that this assumption is unwarranted, and instead, we argue that learning by imitation is actually complex and cognitively challenging.

For instance, Katona [15] explored a matchstick problem with three groups: a creative group that was provided hints for solving the problem, a memory group that saw demonstrations, and a control group that had no help. The creative group performed the best on the experimental task, and the memory group outperformed the control group. A striking result that attracted little attention was that the memory group, which arguably had an opportunity for learning by imitation, outperformed the control group even on novel problems. This outcome suggests that with a demonstration of even one problem to imitate, people may be able to generalize those strategies to new problems.

Furthermore, heuristics, hints and analogies share many of the problems of explicit learning, namely the need for availability of a mentor, a significant time and effort investment on the behalf of the mentor, and the need of a task to which hints, heuristics or analogies can be applied. Also, Nisbett and Wilson [19] found that subjects are typically not aware that hints are given, and they are not accurate in determining which hints (among real and false ones) are useful. Gick and Holyoak [10, 11] found that people often fail to make use of potentially useful analogies unless their relevance is explicitly pointed out to them.

2. Project Description

In this research, the effect of learning by imitation, reinforcement learning and explicit learning on problem solving performance was compared. To the best of our knowledge, learning by imitation has never been explicitly studied in adult humans in the context of higher-level cognition, such as problem solving tasks, with the goal of understanding its underlying mechanisms. Katona's work aimed to show the superiority of hints over demonstrations (memory group). He did not compare the memory group with an instruction group to control for the amount of information subjects got, and he was not interested in studying the mechanisms underlying learning by imitation.

We limited our study to so-called *well-structured* problems [22]. Such problems are characterized by their clear initial and goal states, and by their precisely defined operators and constraints. Furthermore, this research focused on planning-intensive tasks. The Towers of Hanoi problem is a classical example of a well-structured, planning-intensive task.

For this research, a well-known mathematical problem, the ball-weighing problem, was selected. This class of problems has been previously used in a psychological

experiment, the Coin problem [21]. It can be described as follows: "Suppose you have eight coins and a balance. One of the coins is counterfeit, and therefore is lighter than the others. How can you find the counterfeit coin by using the balance only twice?" [5] We used the following variant of the ball-weighing problem: (1) the target ball can be heavier or lighter, (2) 12 balls are used, and (3) the scale can be used 3 times.

The reason for selecting a relatively difficult problem was to require learning, and thus enable differential performance between the experimental groups. Simple problems might not enable the various learning techniques to effectively show how they vary in efficiency.

3. Experimental Design

This section presents the experimental design for studying learning on the ball-weighing problem.

3.1. Design Variables

There were two independent variables in this design. The first one, the experimental *Group*, was a between-subject factor with three levels:

1. Imitation learning group – had access to 5 successful demonstrations of how to solve the target problem (for different ball/weight combinations)
2. Explicit learning group – had access to verbal instructions for solving the problem that they could study for 10 minutes.
3. Reinforcement learning group – got feedback on their performance (whether their answers were correct or not).

Note that the imitation learning and explicit learning groups did not receive any feedback.

The second variable, called trial *Quartiles*, was a within-subject factor. A trial is a single problem instance from its initial presentation until the answer is given. Each trial had different target ball and weight selected at random among the 24 possibilities (12 balls x 2 weights {heavy, light}). Trials were clustered in four quartiles to mark the progression of time within the problem solving session. Besides accommodating the unequal number of completed trials between subjects, this clustering allowed the study of dynamic effects, i.e., how dependent variables evolved over trials.

The design had two dependent variables: elapsed time and accuracy (i.e., whether the answer was correct or not). Both dependent variables were measured on each trial and averaged over trials within each quartile.

3.2. Experimental Hypotheses

Two experimental hypotheses were tested in this experiment.

First, the imitation learning and the explicit learning groups were expected to outperform the reinforcement learning group, both in terms of accuracy (i.e., higher correct answer rate) and speed (i.e., shorter elapsed time per trial). The imitation and explicit learning groups got full information on exactly what to do; in machine learning terms this was supervised learning using fully specified target vectors. In contrast, the reinforcement group only got a binary signal indicating whether the answer was correct or not. In machine learning, reinforcement learning is considered difficult “because the agent is never told what the right actions are, nor which rewards are due to which actions” [26]. Therefore, machine learning predicts that learning by demonstration or supervised learning would outperform learning by reinforcement.

Furthermore, assuming that the verbal instructions were understandable, no difference between the imitation and explicit learning groups were expected because the amount of information given was identical, although presented in a different form.

Second, some kind of learning effect was expected in all groups. Participants were expected to get both more accurate and faster with practice. This effect was expected to be largest in the reinforcement learning group because, in the absence of any information on correct solutions, there was more opportunity to explore wide ranges of solutions of various efficiencies while using feedback to reinforce the better ones. Because such exploration initially requires more processing (i.e., thought), it was expected that the average time to solve the first trials would be longer than the average time to solve the last trials.

3.3. Methods and Procedures

Participants were McGill undergraduate and graduate students. The 68 participants tested have yielded 63 (17 males and 46 females) usable data samples (21 per experimental group). Participants were excluded because they could not finish the warm-up task within 30 minutes (n=3), or because they were identified as statistical outliers on a q-q plot graph (n=2). Participants were randomly assigned to experimental groups. The chance to win a \$50 prize encouraged maximal performance by keeping participants motivated.

A warm-up task (level 1) was first presented (3 balls, 2 uses of scale) to allow participants to become familiar with the task and the user interface.

The target task (12 balls and 3 uses of the scale) was presented as level 2. Upon entering that level, participants were presented with appropriate demonstrations or instructions depending on the condition. They then worked on problem trials for 30 minutes, or until they successfully solved all 24 different trials consecutively. Trials were selected in random order from a list of unsolved trials. When an error was made, the list was reset back to

the whole 24 trials set. Participants got a different trial each time, regardless of whether they successfully resolved the previous one or not.

Participants were instructed to label (categorize) balls to reflect the information they gathered about the balls as their problem solving effort unfolded. Each ball could be labelled as follows: Unknown (heavy, light or normal weight), heavy or light weight, heavy or normal weight, light or normal weight, heavy weight, light weight, or normal weight.

Figure 1 shows a screenshot of the Java computer program designed to implement the ball weighing task and record problem solving data for further analysis. Figure 2 presents the screenshot of the complete instruction set used for the explicit learning group.

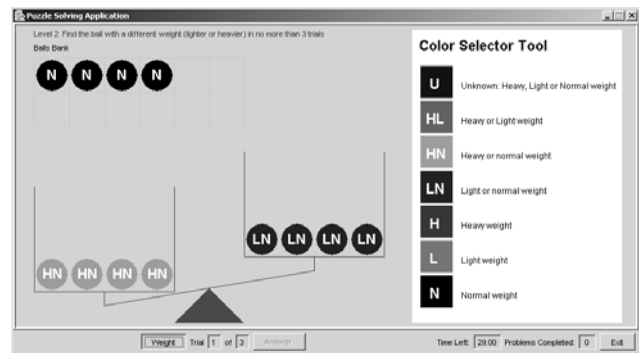


Figure 1. Computer Program Screenshot

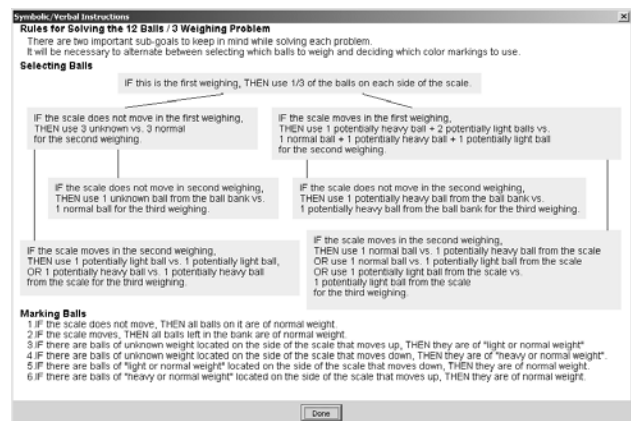


Figure 2. Complete set of explicit instructions

Each participant in the explicit learning group saw a subset of these instructions designed to match the information embedded in a set of five randomly selected demonstrations presented to the imitation learning group.

In terms of task analysis, the ultimate goal of this problem is to identify the target ball, which is either heavier or lighter. To do this, the problem solver must loop through a pair of sub-goals until the problem solver finds a solution or has no further uses of the scale. The first sub-goal

is to select which balls to weigh in order to maximize the information obtained from the scale. The second sub-goal is to appropriately extract new information acquired using the weight trial and to update the state of the balls (i.e., categorize) with the color markings accordingly. Figure 2 presents detailed operations to be performed for all sub-tasks. Note that operations to perform generally depend on the result of the previous weighing.

4. Results

Tables 1 and 2 respectively present mean correct answer rates and elapsed times. Values were computed by averaging trial values for all participants and all quartiles.

Table 1. Mean correct rate¹

<i>Group</i>	<i>Correct Rate</i>	<i>Std Deviation</i>
Reinforcement learning	0.59	0.49
Imitation learning	0.76	0.42
Explicit learning	0.71	0.45

Table 2. Mean elapsed times (ms per trial)

<i>Group</i>	<i>Elapsed time</i>	<i>Std Deviation</i>
Reinforcement learning	104 237	79 020
Imitation learning	87 990	51 068
Explicit learning	94 545	61 180

Contingency tables 3 and 4 present the number of perfect performers and number of correct and incorrect answers for each group respectively. A perfect performer was defined as a participant who made no errors, and thus completed level 2 in exactly 24 trials.

Table 3. Numbers of perfect and imperfect performers

<i>Group</i>	<i>Subjects</i>	<i>Perfect Score</i>	<i>Imperfect score</i>
Reinforcement learning	21	0	21
Imitation learning	21	5	16
Explicit learning	21	2	19

Table 4. Numbers of correct and incorrect answers

<i>Group</i>	<i>Trials count</i>	<i>Correct</i>	<i>Incorrect</i>
Reinforcement learning	350	207	143
Imitation learning	209	311	98
Explicit learning	390	276	114

¹ Computed as follows: Total number of correct answers / Total number of completed trials

5. Analysis and Discussion

5.1. Correct answer rates

Figure 3 presents the mean correct answer rate per group and quartile. The performance decrease in the last quartile might suggest a fatigue effect. However, this effect is not reliable. When trials are grouped into different numbers of clusters, the shape of the distributions varies, as illustrated in figure 4, which presents the mean correct answer rates across 8 clusters. In figures 3 and 4, reinforcement, imitation and explicit learning groups are labeled as RL, IL and EL respectively.

Correct answer rates were not normally distributed and could not be transformed to a normal distribution because of a ceiling value at 1.00, therefore, non-parametric tests had to be performed to determine the statistical significance of the group differences in means. Kruskal-Wallis and Median tests were applied. Separate analyses had to be done for the effect of group and the effect of quartile because those tests allow only one independent variable to be tested at a time.




Figure 3. Mean correct rate per trial quartile (i.e., 4 bins) (n=63 participants)




Figure 4. Mean correct rate for 8 bins (n=59 participants²)

The following analyses were performed using quartiles, i.e. 4 clusters.

5.1.1. Main effect of Group

Kruskal-Wallis and Median tests showed significant differences in correct answer rate across *Group* (Chi-square=7.054, df=2, p=0.029* and Chi-square=7.255, df=2, Median=0.667, p=0.027*, respectively). Pair wise Kruskal-Wallis tests were performed to determine which differences were significant. The results are the following:

1. Reinforcement vs. Imitation learning groups: Chi-square=5.368, df=1, p=0.021*
2. Reinforcement vs. Explicit Learning groups: Chi-square=4.793, df=1, p=0.029*
3. Imitation vs. Explicit Learning groups: Chi-square=0.430, df=1, p=0.512

5.1.2. Main effect of Quartile

The main effect of *Quartile* did not turn out significant under the Kruskal-Wallis test (Chi-square=3.443, df=3, p=0.328) nor under the Median test (Chi-square=6.43, df=3, p=0.092).

5.1.3. Conclusion on correct answer rate

Two conclusions can be drawn. First, the reinforcement learning group significantly underperformed compared to the other two groups, whereas the latter did not differ significantly. Second, a visual inspection of correct answer rate graphs with sufficient numbers of clusters

² Four participants were excluded because they completed fewer than 8 trials.

(e.g., figure 4) suggests a possible modest learning effect, which fell below the statistical power available in this experiment perhaps due to an insufficient sample size.

5.2. Elapsed time

An ANOVA test was performed across the two independent variables (*Group* (3 levels) and *Quartile* (4 levels)) after a log transformation was applied to achieve better normality in elapsed times distributions.

There was no significant main effect of *Group* ($F_{2,60} = 1.42, p = 0.25$). However, the main effect of *Quartile* was highly significant ($F_{3,180} = 85.7, p < 0.001$) suggesting a learning effect across all groups. Furthermore, the interaction effect of *Group* and *Quartile* was significant ($F_{6,180} = 2.63, p = 0.018$) indicating that groups differ on their decrease in elapsed time across *Quartile*. These results suggest different rates of learning, and inspection of figure 5 suggests that the explicit learning group gets the highest speedup. It may have taken time and practice for participants in the explicit learning group to figure out how to effectively make use of the abstract instructions they were given.

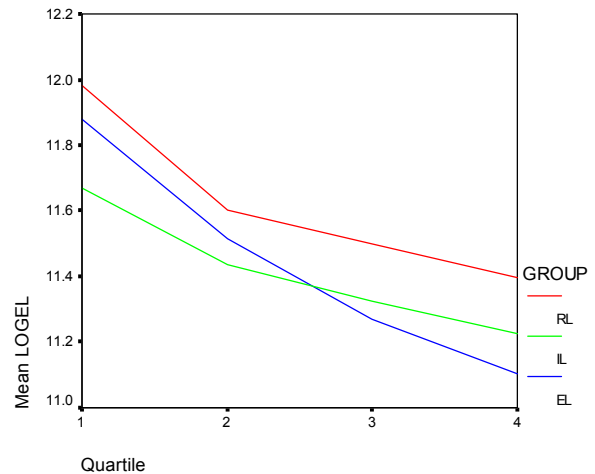


Figure 5. Log of average elapsed time per trial quartile

5.3. Contingency tables analysis

The two Chi-square tests performed on contingency tables of correct answer data were both significant (see Tables 3 and 4.). In other words, *Group* had a significant effect on the number of perfect performers (Chi-square=6.11, df=2, p<0.05) and on the overall number of correct and incorrect answers (Chi-square=21.08, df=2, p<0.05). Most of perfect performers and correct answers are found in the Imitation Learning group.

5.4. Analysis of Strategies

The use of strategies for selecting balls in the first weighing was investigated. The correct strategy is to weigh four balls on the right side of the scale against four balls on the left side (abbreviated as 4/4 below). Besides the 4/4 strategy, two other ones were frequently used: 6/6 (six balls on each side of the scale) and 3/3. Table 5 presents the use of each strategy in each group at the beginning (init) and the end (final) of the 30-minute problem solving session. A strategy was tagged as *n/n* only when it was consistently and exclusively used during the first 20% of trials (init) or the last 20% of trials (final). All other cases (e.g., 2/2 or use of multiple strategies) were categorized as Other.

Table 5. Use of initial and final strategies during the problem solving session for all experimental groups

Strategy	Experimental Group					
	Reinforcement learning		Imitation Learning		Explicit Learning	
	Init	Final	Init	Final	Init	Final
3/3	0.229	0.210	0.0	0.0	0.103	0.167
4/4	0.315	0.623	1.0	0.992	0.833	0.833
6/6	0.349	0.119	0.0	0.008	0.032	0.0
Other	0.107	0.048	0.0	0.0	0.032	0.0

As Table 5 exhibits, participants in the imitation learning group consistently used the correct strategy (4/4) all through the problem solving session.

Participants in the explicit learning group mainly used the correct strategy from the beginning (about 83%), but also explored other possibilities, suggesting it was difficult for them to map the abstract verbal description into action. Results also suggest that a significant proportion of participants in that group interpreted “use 1/3 of the balls on each side of the scale” as suggesting a 3/3 strategy instead of 4/4. The effect was persistent until then end: about 17% of subjects were still using the wrong strategy. This suggests that verbal rules are much more difficult to turn into correct action than a demonstration is, and that although the rules are properly written, they might be misinterpreted.

Finally, participants in the reinforcement learning group increased their use of correct strategy by 31%, indicating a learning effect. The 6/6 strategy was the most popular initially, possibly showing a natural (yet incorrect) heuristic bias that testing all the balls gives the most information. The fact that uses of incorrect strategies (3/3 and 6/6) still remains high in the reinforcement learning group suggests evidence for some kind of block effect that makes it difficult for participants to abandon the use of incorrect strategies. In fact, even incorrect strategies are rewarded because they do yield some correct answers.

Many sub-optimal strategies leave two possible answers at the end, which means 50% of correct answers. If participants simply attempt to “satisfice” [23] (e.g., get 50 or 60% of correct answers), they may well remain stuck with a sub-optimal strategy. Only an optimal solution strategy, such as the one outlined in Figure 2, will always yield the correct answer.

6. Conclusions

Based on the non-parametric and ANOVA statistical test performed on accuracy and speed, the first prediction made has been partially confirmed: the imitation and the explicit learning groups both outperformed the reinforcement learning group in terms of accuracy. However, the other part of the hypothesis was not supported by the data: no significant overall difference in speed (response time) was observed across groups.

The difference in accuracy shows that demonstrations and instructions were significantly better sources of information than simple binary feedback. The experiment also suggests that demonstrations were more effective than instructions based on the contingency tables analysis of correct answers and perfect performers. However, the difference was not large enough to yield a significant under the Kruskal-Wallis and the Median tests. Furthermore, recall that learning by imitation can be done at a much lower cost than explicit learning: the mere presence of a skilled demonstrator solving the task was a sufficient source of information. There was no need for an explicit writing of problem solving rules with all the difficulties it implies. Besides the issues discussed previously (finding a skilled problem solver able to express his strategy in an explicit fashion, making sure the rules are complete and consistent, presenting rules in an appropriate format), we actually found that about 15% of participants misinterpreted the instructions given which is definitely an additional pitfall of explicit learning through abstract verbal rules. In short, learning by imitation was the best method among the three both in terms of accuracy and efficiency.

Learning effects were also found in all groups, although we did not find that the reinforcement learning group was generally slower than the other groups, nor that its speed increased more over trials. If anything, it was observed that the explicit learning group’s speed increased most perhaps because it is more difficult and therefore takes more practice for participants receiving information in an abstract form to figure out how to use it.

In short, accuracy was affected most by the type of learning, whereas elapsed time was affected most by the number of trials completed (i.e., quartile) and by quartile in interaction with group.

This work represents the first step towards showing the importance of learning by imitation in problem solving. The next step will consist in determining more precisely

which mechanisms underlie improved performance in problem solving tasks. Possible mechanisms are the rote memorization of the solution, the priming of certain states and operations, and the acquisition of a complex hierarchical problem representation. Individual differences are also possible, i.e., participants might use different learning mechanisms or even combinations of mechanisms. To explore those underlying imitation learning mechanisms, follow-up experiments will be devised. For instance, a group could be presented with a simpler version of the target task (e.g., 9 balls). Because the task demonstrated is not the same as the target task, rote memorizing could be discounted as an explanation for improved performance in this group.

Our ultimate goal is to model human imitation learning in problem solving tasks using neural networks. Besides helping us uncover the mechanisms underlying learning by imitation, this experiment was used to gather exhaustive data on human problem solving (i.e., each problem solving step is recorded) to train neural networks. Together, experiments and neural network simulations can help devise precise computational models of imitation learning.

References

- [1] A. Alissandrakis, Nehaniv, C. L. and Dautenhahn, K. "Learning by Experience from Others – Social Learning and Imitation in Animals and Robots", *Adaptivity and Learning: An interdisciplinary debate*, Edited by Kuhn, Menzel, Menzel, Ratsch, Richter and Stamatescu, Springer, Berlin, Germany, 2003.
- [2] Anderson, J. R. *Learning and memory, an integrated approach. 2nd ed.*, John Wiley & Sons Inc., New York, NY, 2000.
- [3] G. Ashby and Ell, S. "The neurobiology of human category learning", *Trends in Cognitive Sciences*, Vol. 5, No. 5, May 2001, pp. 204-210.
- [4] Bandura, A. *Social learning theory*, Prentice-Hall, Englewood Cliffs, NJ, 1977.
- [5] Benjafeld, J. G. *Cognition, 2nd edition*, Prentice-Hall, Englewood Cliffs, NJ, 1997.
- [6] C. Breazeal and Scassellati, B. "Robots that imitate humans", *Trends in Cognitive Sciences*, Vol. 6, No. 11, 2002, pp. 481-487.
- [7] R. W. Byrne and Russon, A. E. "Learning by imitation: a hierarchical approach", *Behavioral and Brain Sciences*, Vol. 21, 1998, pp. 667-721.
- [8] M. Carpenter, Call, J., and Tomasello, M. "Understanding 'prior intentions' enables two-year-olds to imitatively learn a complex task", *Child Development*, Vol. 73, 2002, pp. 1431-1441.
- [9] G. Gergely, "Rational imitation in preverbal infants", *Nature*, Vol. 415, 2002, p. 755.
- [10] M. L. Gick and Holyoak, K. J. "Analogical problem solving", *Cognitive Psychology*, Vol. 12, 1980, pp. 306-355.
- [11] M. L. Gick and Holyoak, K. J. "Schema induction and analogical transfer", *Cognitive Psychology*, Vol. 15, 1983, pp. 1-38.
- [12] K. J. Holyoak. "Problem solving". Chapter 8 in *Thinking, an invitation to cognitive science 2nd edition, Vol. 3*. Edited by Smith and Osherson, MIT Press, Cambridge, MA, 1995.
- [13] K. J. Holyoak and Thagard, P. "Analogical mapping by constraint satisfaction", *Cognitive Science*, Vol. 13, 1989, pp. 295-355.
- [14] C. Holroyd and Coles, M. "The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity", *Psychological Review*, Vol. 109, No. 4, 2002, pp. 679-709.
- [15] Katona, G. *Organizing and memorizing*. Columbia University Press, New York, NY, 1940.
- [16] Lieberman, D. A. *Learning: Behavior and Cognition. 3rd ed*, Wadsworth Thomson Learning, Belmont, CA, 2000.
- [17] Lock, A. and Colombo, M. "Cognitive abilities in a comparative perspective", pp. 596-643 in *Handbook of human symbolic evolution*, Edited by Lock and Peters, 1996.
- [18] Meltzoff, A. and Prinz, W. *The imitative mind: development, evolution and brain bases*. Cambridge University Press, Cambridge, UK, 2002.
- [19] R. Nisbett and Wilson, T. "Telling more than we can know: Verbal reports on mental processes", *Psychological Review*, Vol. 84, No. 3, 1977, pp. 231-259.
- [20] Polya, G. *How to solve it 2nd ed.*, Princeton University Press, Princeton, NJ, 1957.
- [21] M. L. Simmel. "The coin problem: a study in thinking", *American Journal of Psychology*, Vol. 66, 1953, pp. 229-241.
- [22] Simon, H. A. "The structure of ill-structured problems", *Artificial Intelligence*, Vol. 4, 1973, pp. 181-202.
- [23] Simon, H. A. *Models of Man*. Wiley, New York, NY, 1957.
- [24] Shultz, T. R. *Computational developmental psychology*, MIT Press, Cambridge, MA, 2003.
- [25] Solso, R. *Cognitive Psychology, 4th edition*, Allyn and Bacon, Boston, MA, 1995.
- [26] Russell, S. and Norvig, P. *Artificial Intelligence, A Modern Approach*, Prentice Hall, Englewood Cliffs, NJ, 1995.
- [27] Tomasello, M. *The cultural origins of human cognition*, Harvard University Press, Cambridge, MA, 1999.