

A Cultural Evolutionary Model for Artifact Capabilities

Felicitas Mokom¹ and Ziad Kobti¹

¹University of Windsor, Windsor, ON Canada N9B-3P4
mokom@uwindsor.ca and kobti@uwindsor.ca

Abstract

The use of tools or artifacts is essential to the human race and has been the subject of recent research in Artificial Intelligence. How individual agents acquire these capabilities and how they evolve can be considered vital steps towards understanding complex group capabilities. In a previous study, we designed and implemented an extended version of a theoretical model for artifact capability that accommodated biological evolution and learning via exploratory methods. Historical knowledge and genetic algorithms were combined with learning techniques to build agents that could learn either individually from observations of their own behaviour or socially by observation from a distance. In this study, we incorporate a collaborative form of cultural learning into the model in an effort to enhance the artifact capability-learning agents. This is accomplished via the design of a cultural evolutionary model that utilizes genetic and cultural algorithms to complement the cognitive abilities of the agents. Learning agents belonging to a social network cooperate with and benefit from each other by sharing individual experiences. Results obtained from the multi-agent simulation implementation confirm the efficiency of social learning over individual learning and demonstrate the benefits of cultural over biological evolution. They also suggest that as artifacts get more complex, social agents learning via cultural influence outperform those learning by observation from a distance.

Introduction

The ability of humans to learn tool or artifact use, evolve these capabilities and transfer the knowledge to others has been of much interest to various researchers particularly in the cognitive sciences. Archaeologists (Plummer, 2004) are fascinated by the earliest recordings of tool use, philosophers (Preston, 1998) theorize on the importance of tool use relative to human intelligence and behavioural geneticists (Bacher et al., 2010) present arguments on the role of genetics in tool use behaviour. Preston contends in her work that the study of tool use be considered as important as the study of language because it is indicative of the high level cognition that humans are capable of. According to (Petroski, 1992) artifact evolution is driven by functionality rather than failure. Artifacts do not necessarily evolve because they fail at what they were intended for, but rather because they can

always be improved. These improvements are often identified during use of the artifact. Humans use tools by themselves but often combine their tool capabilities. In order to successfully model these complex group capabilities it is essential to understand how humans acquire individual capabilities and how these capabilities change over time.

In this study the terms tools and artifacts are used interchangeably and include any physical object in the environment that a human agent can use towards achieving a goal. The human agent is a rational agent that acts in its best interests, has beliefs about the world, and chooses its actions accordingly (Wooldridge, 2000). Based on the Belief-Desire-Intention (BDI) theory of (Bratman, 1987) the rational agent has beliefs, desires and intentions. The agent's beliefs describe its informative state about the world. Its desires represent what the agent would like to accomplish and are used to devise goals. Its intentions are adopted goals that the agent uses to generate plans or actions that it performs. According to (Acay et al., 2008) tool capability resides within the intentions of an agent and represents plans that the agent can realize with the help of a tool. If an agent has capability for a tool then it has at least one plan that specifies one way to use the tool towards one or more of its adopted goals.

In a previous study (Mokom and Kobti, 2011) we implemented an extended version of Acay *et al.*'s theoretical model for tool capability incorporating biological evolution and learning through exploratory methods. A representation of artifacts and the cognition of an agent that can learn artifact capabilities were provided. Learning techniques from (Russell and Norvig, 1995) were combined with genetic algorithms (GA) to build a multi-agent simulation that evaluated individual and social learning in the form of observational learning from a distance. The social learning agent observed another agent successfully apply an artifact capability without the acting agents' knowledge, noted partial information and subsequently formed a learning goal to apply the same capability.

One limitation of the previously implemented social learning agent is the fact that there must exist another agent in its vicinity that already possesses the capability to use

the tool. This limitation coupled with the contention by (Reynolds, 1997) that cultural evolution evolves faster than biological evolution is the inspiration for the work in this paper. We design and implement a cultural evolutionary model that supports agents that can socially learn an artifact capability without any prior knowledge. This is accomplished via the integration of a GA and a cultural algorithm (CA) with the framework of an artifact-capability learning agent. A learning agent can benefit from being part of a social network where individual experiences are shared by using the experiences of others to enhance its own learning process. Our objectives are to demonstrate how agents can collaboratively learn an artifact capability over time and compare the results to those obtained for observational learning agents.

The next section provides some background on related work. It is followed by our architecture of artifact capability-learning agents. We then provide details on our implementation and experiments conducted, followed by conclusions deduced and future work.

Background

Artifact Use

The subject of tool use particularly in animals has been explored in various fields. (Wood et al., 2005) provide a good background on this. Much of the underlying work involves the effort to understand how animals explore objects. (Power, 2000) provides some insight into exploratory methods utilized by children and animals when they encounter a new tool. He contends that the exhibited behaviour, which can sometimes be genetically predetermined, is species dependent and very much influenced by culture.

Robotic researchers have also explored the subject of tool use. This has involved the development of object recognition mechanisms in robots (Wood et al., 2005) and the creation of industrial robots programmed for specific tool use (Bluethmann et al., 2003). In an effort to investigate robots learning tool use through exploratory methods, (Stoytchev, 2005) provides a representation of a robot that can attempt various actions with a tool, record and remember the effects. (Schäfer and Bergfeldt, 2007) investigate the emergence of complex tool use behaviours acknowledging that they need to combine their efforts with learning and reasoning by agents in order to obtain more useful results. (Noble and Franks, 2002) simulate various social learning methods for tool use concluding that emulation is sometimes a more effective method of learning than imitation because it promotes exploration. Omitted from their research is an evolutionary aspect to their work.

Cultural Learning

Knowledge among humans is often transmitted through experience and cooperation. According to Tomasello et al. (1993) in cultural learning, integrated patterns of behavior accumulate changes across generations of a social group.

They identify three different manifestations of cultural learning namely imitation, instructed learning and learning by collaboration. Cultural evolution describes the change of culture over time and can be used to study the effects of cultural learning.

(Curran and O’Riordan, 2007) simulate the instructed learning form of cultural learning using a teacher/pupil environment. In their study a GA and a neural network are used to evolve a population where fitter individuals are selected as teachers for the pupils of the next generation. The goal was to demonstrate how cultural learning improves the fitness of a population. (Acerbi and Nolfi, 2007) utilize simulated annealing to incorporate cultural evolution in their comparisons between individual learning and the imitation form of cultural learning concluding that a sequence of both yields the best results. Geared more towards robotics, much of their work involves robotic sensors and body schema. (Reynolds and Peng, 2004) capitalize on the emergence of cultural learning in their CA framework to demonstrate the power of learning and adaptation within cultures.

Architecture of an Artifact Capability-Learning Agent

An artifact capability-learning agent has the ability to employ learning techniques towards acquiring a tool capability, that is one way to use a tool towards the achievement of one or more of its goals. This can be accomplished via individual or social learning experiences. In the former case the agent learns solely through observations of its own behaviour and in the latter, the agent learns by observing or cooperating with other agents in its environment. We present the existing model of an artifact capability-learning agent from our previous study and demonstrate its expansion to cultural learning agents.

Cognitive Elements of Learning Agent

The cognition of a rational agent endowed with the ability to learn artifact capabilities, was incorporated into a general model of learning agents developed by (Russell and Norvig, 1995). This can be appreciated in Figure 1 obtained from (Mokom and Kobti, 2011). The learning agent’s cognition is composed of a performance element, a critic and a learning element. The performance element bears the responsibility of selecting the agent’s external actions to perform. Once these actions are performed, the critic element measures resulting percepts against an external predefined standard of performance and generates feedback. This feedback is received by the learning element and used to improve the performance element so it can do better the next time. The rational agent’s beliefs, goals and capabilities reside within the performance element playing a role in the decision making process of action selection.




Figure 1: An artifact capability-learning agent

Artifact and Agent Model

An artifact is represented as an object made up of one or more parts. Each artifact part is composed of a set of attributes. An artifact attribute has a set of possible values and a visibility property. The visibility property indicates whether an observing agent can copy the value of the attribute chosen by the agent it is observing. Lets consider a *pen* as an artifact. The part *shell* could represent the entire outer layer of the pen with an attribute *hold-position*. The set of values for the hold-position attribute indicate all possible points where the pen can be held. If the hold-position attribute is visible then an observing agent can copy the point at which the pen is held by the acting agent.

An artifact capability-learning agent has beliefs, goals and capabilities. A capability has an abstract functional ability and an ordered list of tasks. Abstract functional abilities represent all the things that an agent can do with an artifact regardless of whether the agent knows how. It is only when the agent acquires the knowledge to use the abstract ability that the agent can be described as having the capability. An agent can therefore select an abstract ability and use it to formulate a learning goal. For simplicity it is assumed in this study that an artifact has a single part and multiple attributes. The ordered lists of tasks represent attribute value sequences that the agent must determine in order to realize the capability.

External to the rational agent is a predefined standard of performance that is goal dependent for every artifact. This standard maintains information about the number of required tasks and the correct attribute value sequences for each task within an artifact capability. The performance standard is used by the critic element in evaluating the results of the agent's actions. The critic's feedback includes an average fitness score for the attempted sequence. The model supports a *range of values* performance standard that provides an inclusive range within which the selected attribute value is constrained to fall. For an attribute value sequence

$V = \langle v_1, \dots, v_n \rangle$ with n attributes, the critic calculates the mean fitness score MF as follows:

$$f(v_i) = \begin{cases} 1, & mn \leq v_i \leq mx \\ \frac{1}{mn-v_i}, & v_i < mn \\ \frac{1}{v_i-mx}, & v_i > mx \end{cases}$$

$$MF(V) = avg \left(\sum_{i=1}^n f(v_i) \right)$$

where mn represents the lower bound of the performance standard for the attribute with value v_i and mx represents its upper bound. The function gives the same score to all values that satisfy the standards' criteria. The rest of the values are scored based on their distance from the required range.

Performance Element

One of the key decisions in the design of a learning agent is the design of the performance element. In accordance with (Russell and Norvig, 1995), the performance element of an artifact capability-learning agent should contain all the information needed by the agent to go about trying to use the tool. This is essentially how the agent deliberates and selects attribute values. We inherit two types of performance elements designed in the previous study and refer to them henceforth as PE1 and PE2. In this study we design a new type of performance element PE3.

All three performance elements' maintain a history of failed attempts in their respective beliefs. A *fitness-based attribute value selection* procedure is used in the selection of attribute values where one randomly chosen attribute of a selected sequence is modified at each attempt. For PE1 and PE2 the selection is based on the fitness of the agent's previous attempts. PE1 supports an agent learning on its own. The agent simply ensures that it does not repeat attribute value sequences that have previously failed. PE2 supports an agent learning socially via observation from a distance. Like PE1, it does not repeat attempted sequences. The variation lies in the fact that PE2 has partial knowledge of the capability at the start of the learning process. In determining new attribute values, the agent only selects and modifies the invisible attributes.

The new performance element PE3 is built to support an agent learning a tool capability through cultural experience. Its selections are based on both the fitness of its previous attempts and the fitness of all other agents that it cooperates with.

Social Network and Cultural Algorithm

PE3 agents that collaborate with each other belong to a social network. In this study, it is not deemed necessary to define a social network with complex relationships. The social

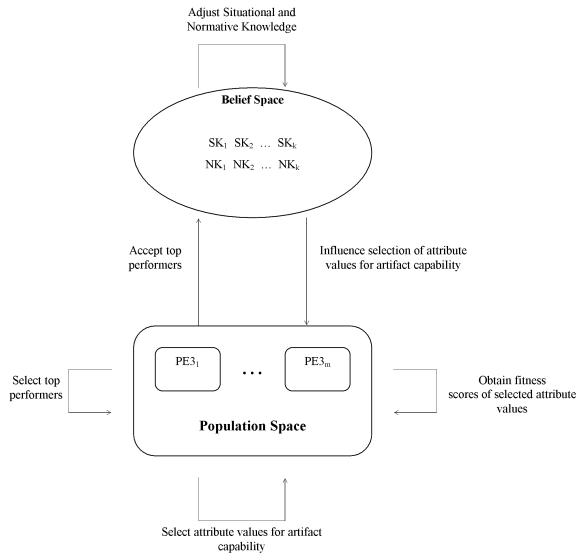


Figure 2: Cultural learning by m agents of a k -task artifact capability

network exists only to facilitate the exchange of information between agents towards enhancing the learning process.

PE3 agents are designed within the CA framework. CA's were introduced by (Reynolds, 1979) to facilitate the modeling of cultural evolution. A CA is made up of a belief space, a population space and a communication protocol between them. Selected individuals from the population space contribute to knowledge maintained in the belief space. The contribution is transmitted through an acceptance function and the knowledge in the belief space is adjusted accordingly. That knowledge influences the evolution of the individuals in the population space via an influence function. A CA supports the use of any kind of evolutionary algorithm in the implementation of the population space. The framework for PE3 agents learning an artifact capability is shown in Figure 2. It demonstrates PE3 agents belonging to a single social network sharing one global belief space within the CA. In the figure, there are m agents trying to learn the same artifact capability with k tasks.

Knowledge Sources Reynolds identifies five types of cultural knowledge that can be maintained in the belief space of a CA. They are situational, normative, topographic, historical or temporal and domain knowledge. Figure 2 shows the belief space in our design using situational and normative knowledge. Situational knowledge maintains the best performers so far. For artifact capability-learning agents cooperating with each other, these would be the highest scoring selections of attribute value sequences. Normative knowledge maintains encouraging ranges for each attribute value making it feasible for agents to “jump into the good range”

(Chung and Reynolds, 1998). The learning agents can adjust their attribute value selections using the guidance of these ranges that have been derived from selections of the top performers.

PE3 agents utilize knowledge from two types of belief spaces. The agent's personal belief space PB maintains a history of its failed attempts for the current task being learned and is local to the agent. Thus $PB = \{\langle v_1, \dots, v_n \rangle\}$ where n represents the number of attributes for the artifact and each v_i is the selected attribute value for the sequence. The global belief space shared by agents in a social network henceforth referred to as GB is defined as: $GB = \langle S, N \rangle$, where $S = \{SK_1, \dots, SK_k\}$ represents the situational knowledge and $N = \{NK_1, \dots, NK_k\}$ represents the normative knowledge for k tasks of an artifact capability. The situational knowledge maintains the single best exemplar found so far for each task:

$$SK = (t, s_{SK}, \langle kv_1, \dots, kv_n \rangle) \quad (1)$$

where t is the task id, n represents the number of attributes for the artifact, each kv_i is the selected attribute value and s_{SK} represents the score of the sequence. The normative knowledge keeps favourable ranges for each attribute value. This is defined as:

$$NK = \{t, R_1, \dots, R_n\} \quad (2)$$

where t is the task id and n is the number of attributes for the artifact. Each R_i is a tuple:

$$R_i = \langle sl, su, [l, u] \rangle \quad (3)$$

where l and u represent the favourable lower and upper bound values of attribute i , with sl and su as their respective scores. A PE3 agent contributes to and uses both belief spaces in the learning process.

Adjusting the Belief Spaces The agent's local belief space is updated with the failed attempt every time the agent tries a new attribute value sequence for a particular task and fails. For GB 's adjustment when top performers are accepted, they are sorted according to their scores. If h contains parameters for the individual with the highest score: $h = (t, s_h, \langle v_1, \dots, v_n \rangle)$, then it is used to adjust the situational knowledge SK defined in Eq. (1) as follows:

$$SK' = \begin{cases} h, & s_h > s_{SK} \\ SK, & otherwise \end{cases} \quad (4)$$

Thus the situational knowledge is always the highest performer so far among all the attempted attribute values for the specific task by members of its social network.

In order to adjust the normative knowledge we need to deal with one attribute at a time. For each attribute i we

obtain and sort its values for all top performers. The lowest selected value x_i and the highest selected value y_i , with their corresponding scores sx_i and sy_i can now easily be extracted. Normative knowledge for the task being learned defined in Eqs. (2) and (3) is updated for each attribute i using the following formulae:

$$\begin{aligned} l'_i &= \begin{cases} x_i, & (x_i < l_i \text{ and } sx_i = sl_i) \text{ or } sx_i > sl_i \\ l_i, & \text{otherwise} \end{cases} \quad (5) \\ sl'_i &= \begin{cases} sx_i, & (x_i < l_i \text{ and } sx_i = sl_i) \text{ or } sx_i > sl_i \\ sl_i, & \text{otherwise} \end{cases} \\ u'_i &= \begin{cases} y_i, & (y_i > u_i \text{ and } sy_i = su_i) \text{ or } sy_i > su_i \\ u_i, & \text{otherwise} \end{cases} \\ su'_i &= \begin{cases} sy_i, & (y_i > u_i \text{ and } sy_i = su_i) \text{ or } sy_i > su_i \\ su_i, & \text{otherwise} \end{cases} \end{aligned}$$

Using these rules, the agents will progress towards learning the correct range required by the performance standard.

Population Space and Influence from Global Belief Space The population space in our cultural algorithm uses a genetic algorithm. As in the previous study the GA uses a bit representation for solutions in the population. It employs two-point crossover and mutation to modify a single attribute value for each attempt. Selection for reproduction is accomplished via roulette wheel selection. With PE3 agents however, mutation is carried out differently. In order to benefit from knowledge in *GB*, the situational and normative knowledge are used to determine direction and step size for the mutation respectively. This effectively permits attribute value sequences to follow the exemplar and at the same time strive to get into a desirable range. If the sequence being influenced is $q = (s, \langle v_1, \dots, v_n \rangle)$, then the chosen attribute's value v_i , is mutated using the following formula derived from Chung and Reynolds (Chung and Reynolds, 1998):

$$v'_i = \begin{cases} v_i + |(u_i - l_i) \cdot N(0, 1)|, & v_i < kv_i \\ v_i - |(u_i - l_i) \cdot N(0, 1)|, & v_i > kv_i \\ v_i + (u_i - l_i) \cdot N(0, 1), & \text{otherwise} \end{cases} \quad (6)$$

where kv_i represents the exemplar value in the situational knowledge as defined in Eq. (1), l_i and u_i correspond to the lower and upper bounds for that attribute in the normative knowledge defined in Eqs. (2) and (3), and $N(0, 1)$ is a random value obtained using the standard normal distribution. All values correspond to the current task being learned.

Cultural Learning Simulation

The simulation environment is a simple 20 x 15 toroidal grid world, in which each square contains an agent and an arti-

fact. There are three types of agents henceforth referred to as AG_GA_PE1, AG_SOCIAL_PE2 and AG_SOCIAL_PE3 varying based on the implementation of the performance element. All AG_SOCIAL_PE3 agents belong to a single social network. The agents can learn capability for artifacts with different complexity. All artifacts are made up of a single part but differ in the number of attributes. The grid is populated with 100 members of each type of agent and the same type of artifact is placed in each square. The agents simultaneously learn the same artifact capability by attempting different combinations of attribute values employing the respective technique of their performance elements.

For the genetic algorithm of AG_GA_PE1 and AG_SOCIAL_PE2 a mutation rate of 0.01 was chosen. For AG_SOCIAL_PE3 agents, mutation was determined by direction and step-size with a mutation rate of $1/n$ where n represented the number of attributes for the artifact being learned. The crossover rate was set to 0.7 and the population size at 100 for the GAs of all agents. The range of possible attribute values for artifacts was set to $[1..100]$ with the *range of values* performance standard covering 20% of the range. The number of tasks required by all agents to learn to use the artifact was 5. Finally, the top 5% performers for each agent's solutions were selected to be accepted into the global belief space.

The pseudo-code for cultural learning of an artifact capability is shown as Algorithm 1. At the start the social network is created and its global belief space is initialized to 5 tasks. For each task the exemplar is set to null and the normative range to the range of possible attribute values that is $[1..100]$. Agents are then added to the network. All agents perform the rest of the algorithm simultaneously. Each agent gets the artifact at its location and uses its cognitive elements to learn the capability. The learning element selects an appropriate ability and formulates a goal. The performance element initializes the agent's local belief and capability to null and the goal to false. Every simulation step the agent generates an attribute value sequence, attempts it and performs the necessary updates. If the feedback of the attempt indicates failure, the sequence is added to the agent's local belief. If there is some form of success, the agent has either reached its goal or has met the minimum requirement to proceed to the next task. If the goal is achieved, the learning element advises the performance element to perform the final capability update and the agent is done. If the goal is not yet reached, the performance element is asked to update the capability with the learned task's successful attribute values, clear the local belief as well as the agent's population space and continue on to the next task.

The knowledge maintained by *GB* is utilized when PE3 provides an attribute value sequence. The pseudo-code is shown as Algorithm 2. POP_SIZE is a constant that specifies the number of attribute value sequences being evolved by the GA used to implement PE3's population space and

Algorithm 1 Cultural learning of an artifact capability

```
Create social network
Initialize global belief space
Add agents to social network
Each agent gets artifact
Learning element selects an ability
Learning element formulates a goal
PE3 initializes local belief
PE3 initializes goal and capability
while goal not achieved

    PE3 provides attribute value sequence
    Critic tests attribute value sequence
    Critic generates feedback
    Learning element generates changes
    PE3 applies changes
end
```

is set to 100 in our experiments. The initial population is randomly generated without repeating sequences that have been attempted already. Once that is complete PE3 provides attribute value sequences by checking if there are still sequences to be attempted, selecting and returning one. If all attribute value sequences in the population have been attempted, fitness scores are used as the criteria to vote top performers for acceptance into *GB*, which is responsible for its own adjustment. A new population is then generated influenced by *GB*'s situational and normative knowledge. One attribute value sequence is selected from the new population and returned for the agent to attempt.

CA's require the evaluation of the entire population space prior to communication with *GB* via the acceptance function. For an artifact capability-learning agent this means that an agent's selections have no impact on *GB* until after all attribute value sequences in the population have been attempted. This is necessary because the agent has to test every generated sequence and obtain its fitness before the top performers can be identified. There is one more instance however, where it would be useful to update *GB*'s knowledge. That would be when the critic element declares success for an agent either at the task or the goal level. This can occur at any time during the evaluation process of the population space. In this specific case, when applying the learning elements suggested changes PE3 requests that the successful attribute value sequence be accepted into *GB*. It does not vote for top performers since it is likely that the entire population has not yet been evaluated. After the successful sequence is accepted *GB* adjusts its situational and normative knowledge using the same rules as when it receives top performers. This allows the agents to benefit from the success of others.

Experiments and Results

AG_GA_PE1 agents are the individual experience-learning agents that strive to acquire an artifact capability on their

Algorithm 2 PE3's algorithm to provide attribute value sequence

```
if size(POP) < POP_SIZE

    values = Generate random value sequence
else
    if attempted all sequences in POP
        Select top performers from POP
        Accept selected performers in GB
        Generate POP' with influence from GB
        values = One value sequence from POP'
    else

        values = One value sequence from POP
    end
end
```




Figure 3: Average Convergence For All Agents Learning Capability for 4, 8, 12 and 16-attribute Artifacts (Visibility of attributes applies only to AG_SOCIAL_PE2 agents)

own. An AG_SOCIAL_PE2 agent learns socially by observing a capable agent from a distance, copying visible attributes and learning the remaining attribute values on its own. AG_SOCIAL_PE3 agents benefit culturally by cooperating with other agents in an effort to enhance their individual learning abilities. Figure 3 shows the results of 100 representatives of each type of agent learning capability for artifacts with 4, 8, 12 and 16 attributes. Figure 4 shows the results of 100 representatives of both types of social agents learning capability for artifacts with 8, 12, 16, 20 and 24 attributes. In all experiments 25% of the attributes were made visible for AG_SOCIAL_PE2. At the end of each test run, the mean convergence times for each type of agent were computed. These are the average number of iterations needed by the agents to learn the artifact capability.




Figure 4: Average Convergence For Social Agents Learning Capability for 8, 12, 16, 20 and 24-attribute Artifacts (Visibility of attributes applies only to AG_SOCIAL_PE2 agents)

It can be observed in Figure 3 that AG_GA_PE1 agents were outperformed by both AG_SOCIAL_PE2 and AG_SOCIAL_PE3 agents in all conducted experiments. The results show an increase in the difference in convergence rate between the individual and social learning agents as the number of attributes increased from 4 to 16 attributes with the individual learning agents needing more time to learn the capability. An interesting observation in Figure 4 is the difference in convergence rate between the two types of social learning agents. AG_SOCIAL_PE2 learn faster than AG_SOCIAL_PE3 agents for 8, 12 and 16 attributes. However at 20 attributes the cultural learning agents outperform those learning via observation from a distance. The trend continues at 24 attributes as AG_SOCIAL_PE3 agents learn even faster.

In our previous study it was demonstrated that learning socially outperforms individual learning therefore, it is no surprise AG_SOCIAL_PE2 agents do better than AG_GA_PE1 agents. The fact that AG_SOCIAL_PE3 agents outperform AG_GA_PE1 agents supports the contention that artifact capability-learning via cultural evolution should proceed at a faster rate than through biological evolution. To understand the results that show agents learning via observation from a distance outperforming their cultural learning counterparts with simpler artifacts or artifacts with fewer attributes it must be remembered that these agents have partial knowledge of the artifact capability upfront. We believe that this partial knowledge gives AG_SOCIAL_PE2 agents a head start in the learning process. AG_SOCIAL_PE3 agents on the other hand begin with no knowledge of the capability and simply use the best of their social group to improve the process over time. According to Reynolds (1997)

knowledge compiled over time and maintained in the global belief space should guide the learning process such that it improves at every trial. As the number of attributes increase, the artifacts get more complex and the search space larger AG_SOCIAL_PE3 agents get better and eventually outperform AG_SOCIAL_PE2 agents. Although the observed threshold may vary and be problem dependent CA's have been used to optimize complex applications Chung and Reynolds (1998). Therefore we suggest that as an artifact gets more complex the likelihood that its capabilities would be best acquired via cultural learning increases especially when the visibility of attributes for observational learning is low.

Conclusions and Future Work

In this study, we have designed and implemented a cultural evolutionary model supporting an agent with the objective of learning artifact capabilities without prior knowledge. Cultural learning agents benefit from belonging to a social network where individual experiences are shared. The model was designed by integrating a genetic and cultural algorithm with the framework of an artifact capability-learning agent. One of the objectives was to enhance the learning capacities of a previously implemented learning agent through cultural learning. Another objective was to compare cultural learning of artifact capabilities to observational learning from a distance. On a larger scale, we maintain that understanding the acquisition and evolution of artifact capabilities for single rational agents is a vital step towards representing their capacity to combine them into group capabilities, towards the accomplishment of more complex goals.

Results obtained from our multi-agent simulation implementation confirm that social learning outperforms individual learning and suggest that complex artifact capabilities are best learned via cultural learning. Although observational learning from a distance surpassed cultural learning for simpler artifacts, the fact that it requires access to an agent already in possession of the capability is a drawback. Additionally the agent must know how to copy the visible attributes with some degree of certainty. A cultural learning agent needs no capable agent in its vicinity and can begin the learning process without possessing any aspect of the artifact capability.

We believe that further experiments are necessary to investigate varying degrees of attribute visibility for agents learning via observation compared to the cultural learning process. One of the knowledge sources identified for the global belief space by (Reynolds, 1979) that was not used in this study is domain knowledge. For future work it can be used to influence the choice of goals for agents to pursue with regard to learning an artifact capability. It would be useful to simulate how goals evolve. One fitness function would no longer be sufficient in the learning process. The choice of a fitness function would be driven by the arti-

fact capability being learned, even for the same artifact. As an example, a knife can be used both as a cooking utensil and as a weapon. An agent's choice of one versus the other would require a different fitness function.

Acknowledgements

This work is made possible by a grant from the National Science Foundation and NSERC Discovery No. 327482.

References

- Acay, D. L., Tildar, G., and Sonenberg, L. (2008). Extending agent capabilities: Tools vs Agents. In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, volume 2, pages 259–265. IEEE Computer Society, Los Alamitos, CA.
- Acerbi, A. and Nolfi, S. (2007). Social learning and cultural evolution in embodied and situated agents. In *Proceedings of the First IEEE Symposium on Artificial Life*, pages 333–340. IEEE Press, Piscataway, NJ.
- Bacher, K., Allen, S., Lindholm, A. K., Bejder, L., and Krützen, M. (2010). Genes or culture: Are mitochondrial genes associated with tool use in bottlenose dolphins (*Tursiops* sp.)? *Behavior Genetics*, 40:706–714.
- Bluethmann, W., Ambrose, R., Diftler, M., Askew, S., Huber, E., Goza, M., Rehnmark, F., Lovchik, C., and Magruder, D. (2003). Robonaut: A robot designed to work with humans in space. *Autonomous Robots*, 14(2):179–197.
- Bratman, M. E. (1987). *Intentions, Plans and Practical Reason*. Harvard University Press, Cambridge, MA.
- Chung, C.-J. and Reynolds, R. G. (1998). CAEP: an evolution-based tool for real-valued function optimization using cultural algorithms. *International Journal on Artificial Intelligence Tools*, 7(3):239–291.
- Curran, D. and O’Riordan, C. (2007). The effects of cultural learning in populations of neural networks. *Artificial Life*, 13(1):45–67.
- Mokom, F. and Kobti, Z. (2011). Evolution of artifact capabilities. To appear in *Proceedings of the IEEE Congress on Evolutionary Computation*.
- Noble, J. and Franks, D. W. (2002). Social learning mechanisms compared in a simple environment. In *Artificial Life VIII*, pages 379–385. MIT Press, Cambridge, MA.
- Petroski, H. (1992). *The evolution of useful things*. Knopf, New York, NY.
- Plummer, T. (2004). Flaked stones and old bones: biological and cultural evolution at the dawn of technology. *American Journal of Physical Anthropology*, 125:118–164.
- Power, T. G. (2000). *Play and Exploration in Children and Animals*. Lawrence Erlbaum Associates, Mahwah, NJ.
- Preston, B. (1998). Cognition and tool use. *Mind and Language*, 13(4):513–547.
- Reynolds, R. G. (1979). *An Adaptive Computer Model of the Evolution of Agriculture for Hunter-gatherers in the Valley of Oaxaca*. PhD thesis, Dept. of Computer Science, University of Michigan.
- Reynolds, R. G. (1997). *Time, Process, and Structured Transformation in Archaeology*, chapter Why Does Cultural Evolution Proceed at a Faster Rate Than Biological Evolution, pages 269–282. Routledge Press, New York, NY.
- Reynolds, R. G. and Peng, B. (2004). Cultural algorithms: Modeling of how cultures learn to solve problems. In *Proceedings of the Sixteenth IEEE International Conference on Tools with Artificial Intelligence*, pages 166–172. IEEE Computer Society, Washington, DC.
- Russell, S. and Norvig, P. (1995). *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ.
- Schäfer, B. and Bergfeldt, N. (2007). Evolution of tool use behavior. In *Proceedings of the First IEEE Symposium on Artificial Life*, pages 31–38. IEEE Press, Piscataway, NJ.
- Stoytchev, A. (2005). Behavior-grounded representation of tool affordances. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3060–3065.
- Tomasello, M., Kruger, A. C., and Ratner, H. H. (1993). Cultural learning. *Behavioral and Brain Sciences*, 16:495–552.
- Wood, A. B., Horton, T. E., and Amant, R. S. (2005). Effective tool use in a habile agent. In Bass, E. J., editor, *Proceedings of the IEEE Systems and Information Engineering Design Symposium*, pages 75–81. IEEE.
- Wooldridge, M. J. (2000). *Reasoning about rational agents*. Intelligent Robots and Autonomous Agents. The MIT Press, Cambridge, MA.