ISSN 2307-4523 (Print & Online)

© Global Society of Scientific Research and Researchers

http://ijcjournal.org/

Learning and Adaptation Strategies for Evolving Artifact Capabilities

Felicitas Mokom*

School of Information Technology, Catholic University Institute of Buea,, P.O. Box 563 Buea, Cameroon, IEEE and IEEE Computational Intelligence Society Member, Association for Computing Machinery (ACM) Member Email: fmokom.dev@gmail.com, fmokom@cuib-cameroon.net

Abstract

In this study we address enhancing the ability of social agents embedded in multi-agent based simulations to achieve their goals by using objects in their environment as artifacts. Reformulated as a discrete optimization problem solved with evolutionary computation methods, social agents are empowered to learn and adapt through observations of their own behavior, others in the environment and their community at large. An implemented case study is provided incorporating the model into the multi-agent simulation of the Village EcoDynamics Project developed to study the early Pueblo Indian settlers from A.D. 600 to 1300. Eliminating the existing presumption that agents automatically know the productivity of the landscape as part of their settling and farming practices, agents use the landscape as an artifact, learning to predict its productivity from a few attributes such as the area's slope and aspect. Given the dynamic nature of the landscape and its inhabitants, agents also evolve various combinations of learning strategies adapting to meet their needs. The result is the demonstration of a mechanism for incorporating artifact use learning and evolution in social simulations, leading to the emergence of favorable learning strategies.

Keywords: multi-agent systems; social simulations; intelligent agents; learning and adaptation.

1. Introduction

Cognitive scientists recognize that tools or artifacts have been vital in the evolution of societies. Archeologists explore early recordings of tool use [1], psychologists study children faced with tool use complexity [2,3,4], behavioral geneticists theorize on the role of genetics in tool use behavior [5], while philosophers argue that tool or artifact use rivals language in depicting the high level cognition attributed to humans [6].

⁻⁻⁻⁻⁻

^{*} Corresponding author.

The relation of artifact use to intelligence has generated an interest in the Artificial Intelligence (AI) community [7] ranging from theories for artifacts and reasoning about their use [8,9], mechanisms for artifact representation and recognition [10,11,12] to models for learning artifact selection and use [13,14,15,16,17,18]. In the realm of autonomous agents and multi-agent systems(MAS) the artifact concept has been introduced as an abstraction to represent the reactive system entities exposed to agents [8,19]. While agents characterize the proactive MAS components responsible for acting upon the environment, artifacts depict the functional components that can be used by agents towards realizing their goals. Essentially agents that can reason adequately about the selection and use of available artifacts can enhance their capabilities in MAS. Omicini and his colleagues artifact theory [8] has been incorporated into the cognition of rational agents formalizing the notion of artifact capabilities [9] used to describe plans that intentional agents can carry out through the use of artifacts. According to the theory the relationship between agents and artifacts can be depicted in three aspects: selection, use and construction. In order to use artifacts towards an objective an agent may need to reason about which artifact to select and once selected how to go about effectively using it for the objective. Artifact construction is often the result of a failure in one of the other two. In this study we focus on artifact use. One of the primary application areas identified to benefit from the artifact abstraction in MAS is multi-agent based simulation(MABS) [19]. MABS has distinguished itself in recent years as one of the most prominent areas of agent-oriented computing providing a technology for simulating complex social systems. In addition to their inherent characteristic of simulating the collective interactions and actions among autonomous and heterogeneous agents, the technology facilitates the natural description of a complex system where agents with the ability to learn and evolve in the presence of others can enter and leave at will [20,21,22]. Assisting in the formulation and validation of theories in the broad field of social science, MABS applications have transitioned from the modeling of simpler societies such as ant colonies [23] to complex human societies such as in the Village EcoDynamics Project (VEP) [24,25,26]. According to Omicini and his colleagues artifacts can be used to describe the function-oriented components of a MABS exploitable by its social agents towards enhancing their capabilities. As such artifacts may encapsulate real world artifacts or any object in the agent's environment that it can utilize. In order to properly reason about using an artifact an agent must have some representation of its structure and possibilities with regard to its behavior. To facilitate this process Omicini and his colleagues introduce the concept of a cognitional artifact in their Artifacts and Agents (A&A) model [19] which they describe as an artifact that exposes three main properties to agents: a usage interface, function descriptions and operating instructions. Function descriptions constitute the services provided by the artifact. Operating instructions for a particular function description are analogous to a manual describing sequences of operations obtained from the usage interface. An agent in possession of a cognitional artifact uses the function description set for artifact selection and a corresponding operating instruction for artifact use. Although cognitional artifacts permit cognitive use of artifacts thus supporting an open MAS, Omicini and his colleagues clearly state that artifacts can still be designed to expose only a subset of the properties to agents. In such cases agents would need to learn and possibly adapt and evolve knowledge for the remaining properties. For MABS systems, the manner in which artifact exploitation evolves can provide even better insight into emerged phenomena and the overall performance its social agents. In order to effectively exploit partially cognitional artifacts in their environments, agents embedded in MABS need a model for artifact use that facilitates learning, adaptation and evolution. Such a model is introduced by [15]. In the study a multi-agent simulation was built composed of agents and artifacts. Artifacts exposed only a usage

interface, defined primitively in terms of functional attributes each with a domain of integer values. The objective of the agents was to learn a correct sequence of actions where each action was a combination of functional attribute values. The cognition of agents in the model was extended from the generic AI model for learning agents[27] to include artifact capabilities. As such the performance element described as the cognitive component used by agents to deliberate on action selection was used to define an individual and a observational learning strategy implemented with genetic algorithms. The authors expanded on their work to include learning and evolving culturally through the use of a cultural algorithm [28]. The studies concluded that while individual learning outperformed random learning, both observational and cultural learning were superior. These studies however fell short by ignoring several relevant MABS components such as the social network connectivity between agents, the dynamic nature of agents, the artifacts and the environment. Nevertheless, they provided an introduction to incorporating artifact use learning and evolution in MABS in a generic way. Based on the artifact capability model [15,29,30], formal definitions of artifacts as exploitable objects in the MABS environment are offered along with agent knowledge structures for reasoning about their use. Agents in our model are built to learn individually, socially through social networks and culturally through a prominent framework for cultural evolution. Agents can use any predefined learning strategy or evolve any combination of these learning strategies over time. To demonstrate the incorporation of our model in MABS the existing multiagent simulation of the Village Ecodynamics Project developed to study the early Pueblo Indian settlers from A.D. 600 to 1300 is utilized as a case study. Agents characterized as households represent and exploit their occupied landscape as an artifact in terms of five functional attribute values with predefined domains. Given the objective of farming for survival, agents employ artifact use learning strategies to predict the productivity of different areas of the landscape directing their decision of where to settle and where to farm. The dynamic nature of the landscape, the mobility of its inhabitants and agents entering and leaving the environment through marriages (new households) and deaths respectively provides a good test bed for the artifact use model. The next section describes the model including artifact and agent representations and the exploitation of artifacts through various learning strategies. It is followed by the case study implementation in the Village simulation. Next, details on conducted experiments are provided, results are discussed, conclusions deduced and some areas for future work are offered.

2. Model for learning and adapting artifact capabilities

To provide a mechanism for incorporating artifact use learning and evolution in MABS, an artifact and agent representation supporting different learning strategies is essential. Artifacts should be defined to expose enough information permitting use to agents that encounter them. Agents should maintain knowledge structures used to reason about artifact use and a mechanism for augmenting their knowledge either on their own or by taking advantage of other agents in their environment. As such of the five cognitive levels identified by Omicini and his colleagues [8] artifact theory ranging from unaware use of artifacts by agents to construction and design of artifacts, this study falls within level three. At this level agents cognitively use artifacts, that is, agents are designed to know which artifact to use but can learn and evolve the knowledge of how to use them over time.

2.1. Artifact Representation

Using the artifact abstraction in the A&A model, artifacts are defined to expose two of the three properties namely function description and usage interface and agents are expected to learn operating instructions. The function description which permits artifact selection is defined to match the agent's objective or an element of the agent's goal set. To facilitate the use of evolutionary computation methods for the learning and evolution process, the usage interface or operations permissible on the artifacts is defined in terms of functional attributes. Hence an artifact is deemed any object that can be described in terms of one or more parts, which have functional attributes with predefined domains[15].

An artifact t is defined as:

$$t \triangleq \langle UI_t, FD_t \rangle \tag{1}$$

 UI_t defines the usage interface as: $UI_t = P_t$ where P_t constitutes the parts of the artifact with each part $p_t \in P_t$ defined as:

$$p_t \triangleq H \, p_t \tag{2}$$

An artifact part is specified in terms of a functional attribute set Hp_t defined as:

$$Hp_t \triangleq \{ \langle [mn,mx], ac \rangle \} \mid mn,mx \in \mathbb{R} \land 0 \le mn \le mx \land ac \in \mathbb{N} \}$$
(3)

where mn, mx specify the lower and upper inclusive boundaries of the domain of its possible values. An additional element of the tuple ac is useful primarily when mn, mx are not whole numbers. It is used to specify the required level of accuracy for the domain elements. In order to learn correct values of an attribute the accuracy is used to convert non-natural real numbers into whole numbers. For example, if the domain of an attribute is defined as [2.115; 6.115] and the accuracy as ac = 2 then the domain is converted to [211; 611] facilitating the discrete optimization process used by the agents in the evolutionary algorithms. FD_t specifies the function description set. A function description $fd_t \in FD_t$ is defined as: $fd_t \triangleq \langle gid , d, x, f \rangle$ where gid is an identifier for the provided service that can be used to match agents goals and d = [0, 1] specifies if the artifact is dynamic with respect to the function description. The number of actions required to successfully exploit the artifact for the objective of gid is denoted as x and f is a fitness function used to evaluate any one of x actions. The fitness function is defined as:

$$f: \langle j, V \rangle \to k \tag{4}$$

where $0 < j \le x$ identifies the index of the action being evaluated and $V = \{v_1, ..., v_n\}$ specifies chosen values for each of the total n functional attributes of the artifact in a maintained sequence. Each value v_i is expected to fall within D_i that is within mn_i and mx_i inclusively. Finally, $k \in \mathbb{R}$ gives a value to the fitness. The dynamic property indicates that the fitness result k may differ during the evolutionary process for the same action index and value set. This is the situation encountered with the exploitation of the landscape in the Village simulation.

2.2. Agent Representation

Mokom and Kobti's [15] agent representation characterizes an agent cognition with a performance element (PE), a critic element (CE) and a learning element. (LE). PE deliberates and chooses the agent's action to perform. Once performed CE evaluates the results obtained against an external performance standard providing feedback to LE which provides any necessary updates to PE. Our agent representation is an extension with an agent *ag* defined as:

$$ag \triangleq \langle PE_{ag}, ce_{ag}, le_{ag} \rangle$$
 (5)

 $PE_{ag} = \langle G_{ag}, C_{ag}, A_{ag} \rangle$ describes the performance element. G_{ag} denotes the agent's set of goals and C_{ag} denotes the agent's set of artifact capabilities. Each goal $g_{ag} \in G_{ag} = \langle gid, st \rangle$ where gid is the goal identifier and st = [0, I] indicates whether the goal is inactive or active. An artifact capability $c_{ag} \in C_{ag}$ is a tuple $c_{ag} = \langle t, CP_{ag} \rangle$ where t denotes an artifact and CP_{ag} a set of plans with $cp_{ag} \in CP_{ag}$ defined as:

$$cp_{ag} \triangleq \langle fd_t, UA_{cp_{ag}} \rangle$$
 (6)

The capability plan cp_{ag} consisting of a function description of artifact *t* and a sequence of use actions is analogous to operating instructions exposed by an artifact in the A&A model. In our model this is maintained within the agent as artifact use knowledge learned and evolved. A use action $ua_{cp_{ag}} \in UA_{cp_{ag}}$ is defined as:

$$ua_{cp_{ag}} \triangleq \langle V, r, y, a_i \rangle \tag{7}$$

where V denotes the functional attribute values pertaining to the action. Agents learning socially may use a predefined social network, build one with a predefined radius or employ a strategy to learn the radius as they evolve. The tuple element $r \in \mathbb{N}$ is used to represent the value of the radius for the action which is only valid when a social learning strategy is employed. The fitness value obtained upon evaluation of the action is denoted by the element $mx \in \mathbb{R}$ and a_i specifies which action generation function or learning strategy the agent used to generate the action. Finally, A_{ag} denotes all the action generation functions available to the agent implementing the learning strategies the agent is able to employ. Essentially, an action generation function $a_j \in A_{ag}$ is defined as $a_j : G_{ag} \to ua$ since the function is used by the agent to generate a use action to perform towards learning a capability that achieves a goal in the agent's goal set. It should be noted that a selected goal will be matched to an artifact via a function description whose usage interface can be exploited as a result. The critic and learning elements of the learning agent are denoted by ce_{ag} and le_{ag} respectively. *ce* which provides feedback from an agent's performed action simply obtains the fitness function result and conveys it to le_{ag} . The final component of the agent model le_{ag} performs any necessary updates on PE_{ag} such as indicating that the goal is achieved or the current action is successful. The performance standard (external to the agent) which learning agents use to evaluate their actions has been incorporated in the artifacts in the form of a fitness function.

2.3. Strategies for learning and adapting artifact capabilities

Given the artifact and agent representations above agents can learn and adapt artifact use by generating actions upon the artifacts towards an active goal. In doing so, they are able to build and ameliorate their capability set for the artifact over time. Agents accomplish this by reformulating the problem as a discrete optimization problem solved with evolutionary computation methods including genetic and cultural algorithms. The functional attribute values and any other values being evolved serve as the discrete variables. The most relevant aspect of the process involves the learning strategies employed by the agents to generate use actions. The model supports three distinct categories of learning strategies within which there are some variations.

2.3.1. Individual learning strategy

The individual learning strategy facilitates agents in MABS learning artifact use through observations of their own behavior. Individual learning is implemented using genetic algorithms (GA) however contrary to the individual learning model in [15] a real-valued representation is used instead of binary for the pool solutions. A solution is a chromosome representing functional attribute values of a use action. To regenerate the pool of solutions after evaluating all its elements, roulette wheel selection is used to select two candidates for reproduction. Two-point crossover is applied at a rate of 0.7 exchanging a single randomly chosen attribute's value. Along with the suggested mutation rate of 1/n for n functional attribute values, real-valued mutation stepsizes are determined using the formula offered by the Breeder Genetic Algorithm [31] with standard values (0.1, 16) for mutation range and mutation precision respectively. The formula generates small step sizes with a higher probability and large step sizes with a lower probability. To address function descriptions characterized by a dynamic property that indicates different results for the same action at different times of the evolutionary process, the fitness score obtained after evaluating a performed action is used to update all pool solutions equivalent to the performed action. As such identical pool elements that have been evaluated always have the same score, that is, the most recent one.

2.3.2. Social learning strategies

One of the benefits for agents embedded in MABS is that they can enhance their learning and evolution with positive influences from other agents in their environment. These influences can be transmitted through an agent's existing social network or one constructed for the sole purpose of artifact use learning. Social networks can remain fixed throughout the model or dynamic in nature where agents continuously update them with new members. Social learning is implemented with a GA similar to the individual learning strategy with a few distinctions. First the chromosome is extended with one more values, representing the radius of the social network utilized by the agent. This value can remain fixed for all agents or given a predefined domain that agents use to evolve its values in the same fashion as the artifact's functional attribute values. Social learning only requires a single-solution pool evolved with the influence of others. The agent searches its network for a better performer whose current action is used to influence the solution. The influence formula to influence solution $W = \{w_1, \ldots, w_{n+1}\}$ with a better performer's solution $Z = \{z_1, \ldots, z_{n+1}\}$ is derived from Chung and Reynolds [32] characterization of influence from an exemplar:

$$w'_{i} = \begin{cases} w_{i} + |(z_{i} - w_{i}) . N(0,1)|, & w_{i} < z_{i} \\ w_{i} - |(z_{i} - w_{i}) . N(0,1)|, & w_{i} > z_{i} \\ w_{i} + (z_{i} - w_{i}) . N(0,1), & otherwise \end{cases}$$
(8)

where *n* represents the number of functional attribute values with one value added to represent the radius and N (0, 1) is a random value obtained using the standard normal distribution. Although agents begin with randomly generated solutions at the start of the simulation, agents that enter the simulation world during the evolutionary process need not commence learning from scratch. In the model, these agents use the latest evaluated use action of their nearest neighbor to begin the learning process.

2.3.3. Cultural learning strategies

In order to simulate cultural evolution the simulated social population in a MABS can be extended to include a cultural evolutionary framework such as a cultural algorithm (CA). The product is a powerful tool for examining and capturing the effects of social interactions and cultural learning on the overall performance of a complex social system [33,14]. Introduced by [34] CAs provide a framework for extracting, storing and exploiting experiences in a population of individuals over time permitting self-adaptation in an evolving model [33]. A population space and a belief space are the two major components in a CA connected via a communication protocol. Selected individuals from the evolving population space which may constitute a social population contribute their experiences to the belief space through an acceptance function. The belief space maintains these experiences as categories of knowledge sources. The types of knowledge sources supported by the CA's belief space are broadly defined to include situational, normative, topographic, historical or temporal, and domain knowledge [35]. The knowledge sources are used selectively or conjunctively to influence the evolution of the individuals in the population space by means of an influence function. The interaction and support that occurs between the population component and the belief space is considered analogous to the evolution of human culture [36]. In this study a CA is utilized to implement the cultural learning strategies. Two types of knowledge sources are maintained in the belief space with the same structure as in Mokom and Kobti [28]. Situational knowledge maintains m best examples extracted from the population where each exemplar is stored as a use action. In the model m can be a fixed value or a percentage of the current population size. Normative knowledge maintains favorable ranges for evolving values specifying the lower and upper boundaries for each functional attribute extracted from the best examples. To facilitate agents that wish to combine cultural learning with social learning the belief space supports maintaining best examples and favorable ranges for social network radius when it is being learned. The belief space is updated at predefined intervals during the evolutionary process. When learning non-dynamic function descriptions, the belief space update uses the same formula in [28] where the current best examples only replace belief space examples with a lower score. For a dynamic function description where the artifact produces different fitness for the same action during evolution the belief space update replaces all examples in the belief space with the current extracted exemplars using them to extract the favorable ranges for the normative knowledge. The population space characterized by the social agents uses a GA similar to the social learning strategy with influence on solutions coming from the belief space rather than better performers of the agent's social network. The situational and normative knowledge sources are used to provide three types of influences: situational, normative and a combination of the two. This extends the CA in [28] which only supports the combined influence. Given solutions W and Z from Section 2.3.2 with Z representing a randomly chosen exemplar situational knowledge influence uses Eq. 8. Normative knowledge influence uses the following formula also derived from Chung and Reynolds [32]:

$$w'_i = w_i + (u_i - l_i) \cdot N(0,1)$$
 (9)

where u_i and l_i denote the upper and lower bounds for the attribute in the normative knowledge respectively. The combined influence uses the same formula from Mokom and Kobti [28] which is Eq. 8 with the intervals $z_i - w_i$ replaced by $u_i - l_i$. Agents are influenced by situational knowledge, normative knowledge or a combination of the two.

2.3.4. Combining strategies

Agents can decide upon any combination of learning strategies to employ. For instance, an agent that wishes to learn on its own as well as socially or culturally would maintain a pool of solutions rather than a single one, using crossover as specified in individual learning and mutation on each solution with influence from better performers in its network or the cultural belief space respectively. An agent can also combine the social and cultural strategies choosing any of the three belief space influence types and randomly alternating its influence between better performers in its network and knowledge from the belief space. Agents influenced by others through their social network are always influenced by the most recent evaluated use action of the influencing agent.

2.3.5. Evolving strategies

While agents can learn with any designated learning strategy our model supports agents that wish to evolve learning strategies as part of the learning process. At a minimum agents that learn which strategy to employ should outperform those that employ learning strategies at random. An individual meta-learning strategy is supported for evolving the strategies the agent is able to employ. A GA is used with a binary representation for the solutions. A binary string of four bits is used to represent a strategy: $[b_1, b_2, b_3, b_4]$. The first two bits b_1 and b_2 represent individual learning and the combined social and social network radius learning respectively. The strategy is applied when the bit is set to '1' and ignored when set to '0'. The last two bits $[b_3, b_4]$ are used to characterize the four possible values for the cultural learning strategies: no cultural learning, situational influence, normative influence and combined situational with normative influence. Disregarding the random strategy [0000] agents learn to evolve a total of 15 learning strategies. The population size is set to 8. Roulette wheel selection is used for selecting solutions for reproduction. Crossover occurs at a rate of 0.7 with two-point crossover applied to alter a single type of influence, that is, b_1 , $[b_2]$ or $[b_3, b_4]$ and a mutation rate of 0.1 is used. Agents that evolve strategies are equipped with the GA for evolving strategies and a separate evolutionary algorithm (GA or CA) for each possible strategy. When evolving strategies, the agent first generates a strategy action which identifies the strategy it wishes to use. The selected strategy is then matched to its evolutionary algorithm which is then used to generate the use action. Since use actions maintain information about the strategy used for their generation, strategies can be evolved alongside the agent's attribute value selections.

3. Case study: Artifacts in the village multi-agent simulation

In this section a case study implementation of the artifact use model for MABS is presented. Involving researchers from several disciplines including anthropology, geology, economics and computer science the

Village Ecodynamics Project was developed to study the early Pueblo Indian settlers from A.D. 600to 1300. A major component of the project, the Village multi agent based simulation, developed using the Repast Simulation Toolkit, models the families within households represented as agents as they farm for maize, hunt for protein, gather water and wood and employ various exchange models for trade [37,24,25]. In the simulation, agents give birth, get married and form new households and die from natural causes or failure in meeting their needs for survival. Soil productivity, rainfall, forest density and animal density are among many other aspects included in the model. Each household carries the responsibility of maintaining plans adaptable to changes in the environment in order to survive. The project strives to understand what led to the depopulation that occurred at the end of period, demography, settlement distributions and violence. The focus in this study is on the farming task conducted by the agents. In the existing model the landscape is divided into cells and agents are presumed to automatically know the soil productivity of every cell at any time that they occupy it. As such agents choose the more productive areas to settle and farm upon. This is highly unrealistic as it assumes that agents can adequately predict soil productivity prior to usage over time. A time step in the simulation is a year characterized by four seasons: spring, summer, fall and winter. Agents consume maize throughout the year but are programmed to plant in the spring and harvest in the fall. Agents self-evaluate and will plant additional plots or move when necessary. Many factors are utilized to measure soil productivity which changes over time and declines depending on how long and how often it has been cultivated. Representing the landscape as an artifact (LANDSCAPE), five features that agents are likely to know upon encountering a cell are selected: the average elevation (dem), the average slope(slope), the average direction of slope (aspect), the average depth to bedrock (depth) and the average proportion of its biomass consisting of any subspecies of big sagebrush prior to any agricultural clearing (artr). With supplied predefined domains for each feature, LANDSCAPE is defined as a single part artifact with the five features represented as its functional attributes. The domain of all attributes are converted to integers. The accuracy is set to 0 for all attributes except *artr*, for example *depth* is rounded from [25.2,182.7] to [25,183]. Since artr has a domain of [0.008421053,0.5198181], its accuracy is set to 3 resulting in a conversion to [0,500]. LANDSCAPE has one function description with farming as the provided service requiring one action for exploitation. Since landscape productivity changes overtime resulting in the same functional attribute value combinations producing different results over the years the dynamic attribute is set to 1. Given by the current simulation the fitness of a use action is essentially the agent's harvest. Since agents always plant in their settled cell, agents generate use actions whenever they need to plant in additional cells or move within the predefined move radius. Basically, the agent tries to predict the productivity of included cells. New use actions are only generated when the current one has not received a fitness. Since a combination of functional attribute values in a generated use action will not necessarily be identical to those in a particular cell, an interpretation layer is needed to convert the use action to the closest matching cell. To do this artr values are converted to their integer equivalent and a simple distance measure is used averaging over all attribute values to select the cell that is closest to the generated values. For a given cell with functional attribute values $CV = \{cv_l, cv_l, cv_l,$..., cv_5 and a generated use action with functional attribute values $V = \{v_1, \dots, v_5\}$ the following function is applied:

$$dst(v_i, cv_i) = \begin{cases} 1.01, & v_i = cv_i \\ \frac{1}{|cv_i - v_i|}, & otherwise \end{cases}$$
(10)

 $Dst(V, CV) = avg(\sum_{i=1}^{5} dst(v_i, cv_i))$

If multiple cells have the same distance measure one is randomly selected. Agents in the simulation are extended to support learning to exploit the LANDSCAPE artifact using any of the strategies. They can also evolve strategies. Learning algorithms are applied when the agent decides to move or plant additional plots. First, the agent selects a number of cells for consideration as in the current simulation. Then the agent's PE element determines which strategy to use if evolving strategies. The chosen strategy is then employed to generate a use action. The use action is matched to the closest cell which the agent presumes to be the most productive using Eq. 10. The harvest obtained from the plots farmed is provided as feedback to the agent's CE element which updates the evolutionary algorithms in PE: Over time the agent learns, evolves and adapts as necessary to changes in settlement distribution, landscape productivity, demographics of households and the population at large, its social networks and emergent cultural beliefs. Along the way the agent maintains its primary objective which is to produce enough to feed its family for survival as agents that do not produce enough maize will ultimately die. One important aspect to note about the learning strategy that involves learning the social networks, maintaining such networks in areas with high variability may prove detrimental to the agents.

4. Experiments and results

4.1. Environment setup

The village simulated in conducted experiments is VEP IIN which models a larger region than the original VEP I. Spatially agents occupy a landscape represented as 114,240 cells (VEPI occupied 45,400 cells [25]). Agents are stripped of all tasks except farming and not allowed to any form of trade in order to adequately measure the artifact use effects on the evolution of the population. Agents use only the productivity (or predicted productivity) as criteria for where to relocate and farm. All experiments begin with 600 randomly generated agents randomly placed on the landscape. Each agent is given access to the LANDSCAPE artifact which exposes itself as an object usable for the service of farming. The following domain ranges are given for the functional attributes of each cell: artr [0,500], aspect[0,359], dem [1438,3009], depth and slope [0,49]. Upon arrival on the landscape, agents immediately attempt to move to a predicted productive cell. Obviously at the start of the simulation this is simply a guess as agents have no idea about the fitness of their actions. Over time however it is expected that agents will improve their predictions of cell productivity and therefore enhance their chances for survival. The existing move radius of 40 is maintained and a domain of [1,40] is defined for the social network radius. Since feedback is only obtained once a year after a harvest, population size of the solutions pool for all strategies involving individual learning is set to 10. The contention is that a small population size is relevant to minimize agents dying before learning begins. The population size for the GA used for evolving strategies is set to 8. The belief space maintains the current top 2% of the population and is updated every 5 years. It should be noted that although the number of examples in the belief space begins at 12 with 600 agents in the population, this value changes over time along with the population size. For all experiments the simulation begins in year A.D. 600 and runs through 1280. Results are aggregated over the four study periods obtained from the Pecos classification [38] used by the Village Ecodynamics Project researchers. They are

Basketmaker III (A.D. 600-900), Pueblo I (A.D. 750-900), Pueblo II (A.D. 900-1150) and Pueblo III (A.D. 1150-1280). The objective is to initiate the avenue for validating results against archeological findings.

4.2. Test cases

In the village simulation agents maintain a maizestorage attribute which tracks the maize they have in store at any given time. This attribute is augmented with the agents' harvest and reduced by maize consumed by family members of the agent's household. Factors that affect the amount of maize consumed by an agent include calories spent on farming, monitoring plots and traveling to farm on plots away from its settled cell. Agents die when maizestorage drops below zero. In all conducted experiments, we are interested in the survival rate of the population during all the classified phases. To facilitate identifying the no learning and the various learning strategies employed they are henceforth referred to as follows: No learning (Random), Individual learning (Indv), Social learning with randomly generated radius (SocRRad), social learning with learned radius (SocLRad), cultural learning with situational knowledge (CulS), cultural learning with normative knowledge (CulN) and cultural learning with combined situational and normative knowledge (CulB). In the case of Random the agent generates a use action using randomly generated attribute values. In addition, Original is used to represent the results where agents are assumed to know the soil productivity and *EvStrategy* is used for agents evolving strategies. Results for every conducted experiment shows the number of agents that survived at the end of each classified phase. In the first set of experiments agents learn with specified strategies. The first test case compares the cultural and social learning strategies: CulS, CulN and CulB, SocRRad and SocLRad. Results are shown in Table I. Next we compare agents that are not learning with a strategy from each of the three main categories: Random, Indv, CulB and SocLRad with results shown in Table 2. Finally a comparison is done between combined strategies: Indv+ SocLRad, Indv + CulB, SocLRad + CulB and Indv + SocLRad + CulB. Results are shown in Table 3. The next experiment investigates agents evolving strategies. For this, there is a single test case where the results of agents evolving with randomly chosen strategies are compared to those with learned ones. Results are depicted in Table 4. With agents still stripped of all tasks except farming, trading disallowed and relocation or cells for planting selected only on the basis of productivity, the final experiment compares results for agents that know productivity at all times with the results of SocLRad. For this, the survival of the population is shown in Table 5.

5. Discussion

The conducted experiments examine the survival rate of the population throughout the evolutionary process by tracking the number of agents that survive at the end of each of the classification phases defined in Section 4.1.

Phase	CulS	CulN	CulB	SocRRad	SocLRad
Start	600	600	600	600	600
BasketMaker III	28	85	29	65	303
Pueblo I	6	3	19	31	514
Pueblo II	0	0	0	2	1653
Pueblo III	0	0	0	0	2932

Table 1: Agent survival for cultural-learning and social-learning agents at the end of each classification phase

Table 2: Agent survival for agents not learning compared with agents learning at the end of each classification

phase

Phase	Random	Indv	CulB	SocLRad
Start	600	600	600	600
BasketMaker III	14	21	29	303
Pueblo I	0	0	19	514
Pueblo II	0	0	0	1653
Pueblo III	0	0	0	2932

Table 3: Agent survival for combined learning strategies at the end of each classification phase

Phase	Indv+SocLRad	Indv+SocLRad+CulB	Indv+CulB	SocLRad+CulB
Start	600	600	600	600
BasketMaker III	181	132	42	142
Pueblo I	154	218	1	205
Pueblo II	265	357	0	288
Pueblo III	492	394	0	321

Table 1 shows the results of the test case which compares CulS, CulN and CulB, SocRRad and SocLRad. Although all cultural learning agents die by the end of Pueblo II, CulB agents slightly outlive the others. Interestingly, at the end of Basketmaker III, CulN agents have a higher survival count but are outperformed by CulB agents in the next phase. This raises the question whether a particular strategy can be considered better during a certain phase. Agents learning culturally are influenced by the exemplars in the belief space extracted from the population at large. Results suggest that solutions discovered by the population's best performers may not be appropriate for most agents. Also, the frequency of belief space updates may play a role as agents may start to underperform before a better solution is available. Comparing the cultural learning strategies of SocLRad with SocRRad, it is demonstrated that the size of the agent's network plays a significant role in its chances for survival. SocRRad agents barely survive through Pueblo I with results similar to CulB. SocLRad agents survive all phases with a notable increase in each phase after BasketMaker III. It is likely that these agents learn to use a smaller radius in homogeneous regions and widen their radius in heterogeneous ones. In general, the social learning agents outlive those learning via the cultural learning methods. Table 2 shows very little difference between Random and Indv agents. Neither one makes it to Pueblo I. CulB has a slight edge barely surviving through Pueblo I but once again the superiority of SocLRad is evident. Results with various learning strategy combinations in Table 3 show the importance of social learning on the LANDSCAPE artifact. Although lower than SocLRad employed on its own, the survival count remains positive for any combined strategy that includes it. By contrast Indv + CulB agents are gone by the end of Pueblo I. Once again there some notable performance differences between the strategies in different phases. For instance, while Indv + SocLRad seems outperformed by Indv + SocLRad + CulB in Pueblo I and Pueblo II it ends up with the highest survival count at the end of Pueblo III. Ultimately SocLRad + CulB seems to be the most consistent improving slightly with each phase.

Phase Random EvStrategy Start 600 600 101 BasketMaker III 77 Pueblo I 14 139 Pueblo II 0 154 Pueblo III 0 144

 Table 4: Agent survival at the end of each classification phase for agents randomly choosing strategies

 compared with agents evolving strategies

Results depicted in Table 4 demonstrate that when agents are evolving strategies the population struggles to survive. However, there is a better performance when strategies are evolved rather than not. It is possible that an agent surviving with a particular strategy does not learn an appropriate one in time given the dynamic nature of the environment.

 Table 5: Agent survival for agents in the original simulation that know productivity compared with agents

 learning

Phase	SocLRad	Original
Start	600	600
BasketMaker III	303	2958
Pueblo I	514	9682
Pueblo II	1653	9825
Pueblo III	2932	10338

Table 5 shows the results of the final experiment which compares *SocLRad* to *Original*. Although there is a significant difference in the survival counts, it is evident that *SocLRad* agents survive just as their *Original* counterparts do. Possible reasons for the difference include our representation of LANDSCAPE with only five functional attributes. and the assumption that all agents are completely unaware of the productivity of the soil. If we seeded good solutions into some of the agents at the start such as is done to implement observational learning in Mokom and Kobti[15] it may be possible to achieve results much closer to those obtained by *Original*. In particular since the results demonstrate the effectiveness of knowledge propagation through the social networks a few knowledgeable members of the population maybe enough to significantly influence the survival rate.

6. Conclusions and future work

This study provides a model for incorporating agents learning and adapting the use of objects and artifacts towards realizing their goals. The model extends the artifact abstraction in MAS[19] to include evolution and learning for social and cultural agents. Formal definitions for artifacts and agents reasoning about their use in MABS are provided along with a variety of learning strategies that agents can employ towards enhancing their

capabilities taking advantage of others in their environment. An existing Village multi-agent based simulation [25] which models the lives of the ancient Pueblo Indians spanning almost 700 years is utilized as a case study. With agents represented as households farming for survival the landscape is modeled as an artifact abstraction that agents learn to exploit from a representation of a few of its attributes. Given the fact that the landscape is dynamic with productivity changing over time, and characterized by agents entering, leaving and relocating as they strive to survive, learning and adaptation is essential for every agent. Experiments conducted track the survival of the agents and identify agents learning through social influence concurrently evolving the size of their social network as the most successful strategy. It is important to note that while it proves to be the best with this particular case study different strategies may prove better in other applications. While this model may prove useful in many situations, there are also some limitations. Fundamentally, the model assumes a reduction of an artifact to a set of functional attributes. Since these attributes can be numerous, the model may prove intractable for problems where a small set of attributes would be insufficient to obtain useful results. In addition, as in the case with many agent-based simulations, model parameters often need to be fine-tuned in order to examine their role and impact. For future work it would be useful to compare the findings with the archeological records. Additional attributes could be used to define LANDSCAPE to observe if agents learn better or faster. In enhancing the case study agents could be modeled to exploit real tools for farming or hunting. It may also be useful to investigate the possible effects on survival when a percentage of the agents are aware of the productivity at the start of the simulation. In the model all agents are equally susceptible to positive influences. It may be useful to address the possibility that some agents are resistant. Finally, the model has only measure done aspect, the survival of the agents. It may be relevant to investigate other aspects of the simulation such as family size and spatial distribution of the agents in each phase.

7. Recommendations

To address the fact that a few useful functional attributes for an artifact being studied should be identified in order to obtain useful results from the model, it is recommended that any application of the model to a particular domain be conducted with the assistance of an expert in that domain. This will ensure that the attributes chosen for evolution prove to be relevant for the problem being addressed. These experts can also play a significant role in the tweaking of the model parameters towards the discovery of emergent behavior in the model. This aspect was applied in the case study for the Village Ecodynamics project, where archaeologists selected the features of the landscape that were to be evolved.

Acknowledgements

Collaborations with Dr. Ziad Kobti from the University of Windsor and the Village Ecodynamics project team (Kyle Bocinsky, Stefani Crabtree and Dr. Tim Kohler) from Washington State University.

References

 T. Plummer, "Flaked stones and old bones: biological and cultural evolution at the dawn of technology.," American Journal of Physical Anthropology, vol.125, pp.118–164, 2004.

- [2]. T. G. Power. Play and Exploration in Children and Animals. Mahwah, NJ: Lawrence Erlbaum Associates, 2000.
- [3]. A. K. Gardiner, D. F. Bjorklund, and M. L. G. ad Sarah K. Gray, "Choosing and using tools: Prior experience and task difficulty influence preschoolers' tool-use strategies," Cognitive Development,vol.27, no.3, pp.240–254, 2012.
- [4]. K. Neldner, E. Reindl, C. Tennie, J. Grant, K. Tomaselli, and M. Nielsen, "A cross-cultural investigation of young children's spontaneous invention of tool use behaviours," Royal Society Open Science, vol.7, no.5, p.192240, 2020.
- [5]. K. Bacher, S. Allen, A. K. Lindholm, L. Bejder, and M. Krützen, "Genes or culture: Are mitochondrial genes associated with tool use in bottlenose dolphins (Tursiop sp.)?," Behavior Genetics, vol.40,pp.706–714, 2010.
- [6]. B. Preston, "Cognition and tool use.," Mind and Language, vol.13,no.4, pp.513–547, 1998.
- [7]. A. B. Wood, T. E. Horton, and R. S. Amant, "Effective tool use in a habile agent.," in Proceedings of the IEEE Systems and Information Engineering Design Symposium (E. J. Bass, ed.), pp.75–81, IEEE,2005.
- [8]. A. Omicini, A. Ricci, and M. Viroli, "Agens Faber: Toward a theory of artefacts for MAS.," Electronic Notes in Theoritical Computer Sciences, vol.150, no.3, pp.21–36, 2006.
- [9]. D. L. Acay, G. Tildar, and L. Sonenberg, "Extending agent capabilities: Tools vs Agents," in Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, vol.2, pp.259–265, IEEE Computer Society, Los Alamitos, CA, 2008.
- [10]. E. Bicici and R. S. Amant, "Reasoning about the functionality of tools and physical artifacts," Department of Computer Science, North Carolina State University, Tech Rep, vol.22, 2003.
- [11]. Y. Wu and Y. Demiris, "Learning dynamical representations of tools for tool-use recognition," in IEEE International Conference on Robotics and Biometics, IEEE, 2011.
- [12]. K. P. Tee, J. Li, L. T. P. Chen, K. W. Wan, and G. Ganesh, "Towards emergence of tool use in robots: Automatic tool recognition and use without prior tool learning," in 2018 IEEE International Conference on Robotics and Automation (ICRA), pp.6439–6446, IEEE, 2018.
- [13]. A. Stoytchev, "Behavior-grounded representation of tool affordances.," in Proceedings of the IEEE International Conference on Robotics and Automation, pp.3060–3065, 2005.
- [14]. Z. Kobti, A. Snowdon, S. Rahaman, T. Dunlop, and R. Kent, "A cultural algorithm to guide driver learning in applying child vehicle safety restraint," in Proceedings of the IEEE Congress on Evolutionary Computation, pp.1111–1118, 2006.
- [15]. F. Mokom and Z. Kobti, "Evolution of artifact capabilities," in Proceedings of the IEEE Congress on Evolutionary Computation, (New Orleans, Louisiana), pp.476–483, IEEE Press, June 2011.
- [16]. R. Jain and T. Inamura, "Bayesian learning of tool affordances based on generalization of functional feature to estimate effects of unseen tools," Artificial Life and Robotics, vol.18, no.1-2, pp.95– 103,2013.
- [17]. S. Brown and C. Sammut, "A relational approach to tool-use learning in robots," in Inductive Logic Programming, pp.1–15,Springer, 2013.
- [18]. F. Mokom and Z. Kobti, "Improving artifact selection via agent migration in multi-population cultural

algorithms," in 2014 IEEE Symposium on Swarm Intelligence, pp.1-8, IEEE, 2014.

- [19]. A. Omicini, A. Ricci, and M. Viroli, "Artifacts in the A&A metamodel for multi-agent systems," Autonomous Agents and Multi-Agent Systems, vol.17, no.3, pp.432–456, 2008.
- [20]. E. Bonabeau, "Agent-based modeling: Methods and techniques for simulating human systems," Proceedings of the National Academy of Sciences of the United States of America, vol.99, no. Suppl 3, pp.7280–7287, 2002.
- [21]. N. Gilbert, "Agent-based social simulation: dealing with complexity," The Complex Systems Network of Excellence, vol.9, no.25, pp.1–14, 2004.
- [22]. H. Jiang, W. Karwoski, and T. Z. Ahram, "Applications of agent based simulation for human sociocultural behavior modeling," Work:A Journal of Prevention, Assessment and Rehabilitation, vol.41,pp.2274–2278, 2012.
- [23]. A. Drogoul and J. Ferber, "Multi-agent simulation as a tool for modeling societies: Application to social differentiation in ant colonies," in Artificial Social Systems, pp.2–23, Springer, 1994.
- [24]. Z. Kobti, R. G. Reynolds, and T. Kohler, "A multi-agent simulation using cultural algorithms: the effect of culture on the resilience of social systems," in Proceedings of the IEEE Congress on Evolutionary Computation, vol.3, pp.1988–1995, 2003.
- [25]. T. A. Kohler, C. D. Johnson, M. Varien, S. Ortman, R. Reynolds, Z. Kobti, J. Cowan, K. Kolm, S. Smith, and L. Yap, "Settlement ecodynamics in the prehispanic central mesa verde region," The model-based archaeology of socionatural systems, pp.61–104, 2007.
- [26]. S. A. Crabtree, K. Harris, B. Davies, and I. Romanowska, "Outreaching archaeology with agent-based modeling: Part 3 of 3," Advances in Archaeological Practice, vol.7, no.2, p.194, 2019.
- [27]. S. Russell and P. Norvig. Artificial Intelligence: A Modern Approach. Upper Saddle River, NJ: Prentice Hall, 2010.
- [28]. F. Mokom and Z. Kobti, "A cultural evolutionary model for artifact capabilities," in Proceedings of the European Conference on Artificial Life, (Paris, France), pp.542–549, MIT Press, August 2011.
- [29]. F. Mokom and Z. Kobti, "Exploiting objects as artifacts in multi agent based social simulations," in Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, pp.1869–1870, 2015.
- [30]. F. Mokom, "Modeling the evolution of artifact capabilities in multi agent based simulations,", 2015.
- [31]. D. Schlierkamp-Voosen and H. Mühlenbein, "Predictive models for the breeder genetic algorithm," Evolutionary, vol.1, no.1, pp.25–49,1993.
- [32]. C.-J. Chung and R. G. Reynolds, "CAEP: an evolution-based tool for real-valued function optimization using cultural algorithms," International Journal on Artificial Intelligence Tools, vol.7, no.3,pp.239– 291, 1998.
- [33]. Z. Kobti, R. G. Reynolds, and T. Kohler, "The effect of kinship cooperation learning strategy and culture on the resilience of social systems in the village multi-agent simulation," in Evolutionary Computation, 2004. CEC 2004. Congress on, vol.2, pp.1743–1750,IEEE, 2004.
- [34]. R. G. Reynolds, An Adaptive Computer Model of the Evolution of Agriculture for Hunter-gatherers in the Valley of Oaxaca. Ph.D.thesis, Dept. of Computer Science, University of Michigan, 1979.
- [35]. B. Peng, R. Reynolds, and J. Brewster, "Cultural swarms," in Evolutionary Computation, 2003.

CEC'03. The 2003 Congress on, vol.3, pp.1965–1971, IEEE, 2003.

- [36]. R. G. Reynolds and B. Peng, "Cultural algorithms: Modeling of how cultures learn to solve problems.," in Proceedings of the 16th IEEE International Conference on Tools with Artificial Intelligence, pp.166– 172, IEEE Computer Society, Washington, DC, 2004.
- [37]. T. A. Kohler, "The final 400 years of pre-hispanic agricultural society in the mesa verde region," Kiva, vol.66, pp.191–264, 2000.
- [38]. T. A. Kohler, "News from the northern american southwest: Prehistory on the edge of chaos," Journal of Archeological Research, vol.1, no.4, pp.267–321, 1993.