2014 IEEE Congress on Evolutionary Computation (CEC)
July 6-11, 2014, Beijing, China

Analyzing Prehistoric Hunter Behavior with Cultural Algorithms

Samuel D. Stanley, Areej J. Salaymeh, Thomas J. Palazzolo, and David M. Warnke

Abstract— **This paper details a cultural algorithm (CA) system designed to assist archaeological expedition teams in the task of finding historic artifacts. In our system, the goals that the agents are trying to achieve continuously change as the environment changes. We are thus able to simulate the real-world challenge of a dynamic environment that human cultures must deal with and react to, making our system a very useful tool for finding the archaeological remains of such cultures. Although it is very new, our system has already had yielded promising results in the service of Dr. John O'Shea's Lake Huron expedition team which is studying the prehistoric Alpena-Amberley Land Bridge. We hope to use it to assist other expeditions as well in the near future.**

Index Terms—**Anthropology, Archaeology, Artificial Intelligence, Cultural Algorithms, Huron.**

I. Introduction

THE Alpena-Amberley Ridge, an underwater feature currently submerged under Lake Huron, was actually a dry land corridor during a portion of the Early Holocene which spanned the lake and linked what is now Alpena, Michigan, USA to what is now Amberley, Ontario, Canada. In 2008, Dr. John O'Shea of the University of Michigan hypothesized that when the Ridge was a dry land corridor, it may have provided a dry path for migration of caribou and other animals across Lake Huron and thus a hunting ground for Paleoindian hunter-gatherers. To test his hypothesis, Dr. O'Shea collected underwater samples and data from portions of the Ridge using sonar and underwater autonomous vehicles. He found what appeared to be manmade hunting paraphernalia such as hunting blinds and caribou drive lanes [1]. This find was met with interest from the research community, including the Artificial Intelligence community. This paper details a system that was originally designed to aid Dr. O'Shea's expedition team in finding other lost artifacts in the Alpena-Amberley Ridge region, but can also be used to find artifacts in other environments for the benefit of other archaeological projects.

The first draft of this paper was submitted for review on 1/20/2014 and accepted on 3/15/2014. The revised draft was submitted on 4/15/2014.

S. D. Stanley is a PhD student in the Department of Computer Science at Wayne State University, Detroit, MI, 48202 USA (phone: 313-530-8837; e-mail: sdstanley467@ yahoo.com).

A. J. Salaymeh is a PhD student in the Department of Computer Science at Wayne State University, Detroit, MI, 48202 USA (email: eg6674@wayne.edu).

T. J. Palazzolo is a PhD student in the Department of Computer Science at Wayne State University, Detroit, MI, 48202 USA (phone: 248-797-5160; email: algusunderdunk@gmail.com).

D. M. Warnke is a senior undergraduate with a Major in Computer Science at Wayne State University, Detroit, MI, 48202 USA (e-mail: david.m.warnke@gmail.com).

Fig. 1. The "Dragon Blind", an Interesting Structure Found by Dr. O'Shea's Initial Expedition to the Alpena-Amberley Ridge in 2008 [1]

Given data describing the rate of change of various factors within an environment, our system provides automatic design of the changed environment for all desired intervals within a given time period. This allows researchers to design and perform experiments where environmental change is a crucial factor without having to manually redesign the environment whenever it changes. Additionally, many helpful tools and functionalities for designing and implementing experiments are provided.

When the program is started, the system first generates a 3D simulation world and a 2D hashmap of that world from raw text files containing height information. To provide data for the specific system within the program which is the central focus of this paper, the user also has the option of including text files containing time series data for environmental variables. For all such variables included, the system automatically fills all gaps in their time series data through linear interpolation until there is one value for each year in the period spanned by the series. For the purposes of the main experiments discussed in this paper we have included only a water level environmental variable (although in some other experiments done in the past we also included a temperature variable as well).

The system has two basic modes: The standard mode and the CA-loaded hunting blind finder mode. The first mode allows the user to enter any year and have the system automatically modify the program's 3D simulation world and 2D hashmap so that it accurately represents the historic environment during that year based on the time series data provided. Additionally, it is very easy to run experiments within these 3D historic simulation worlds. For instance, a herd of caribou may be generated and made to migrate from a given start point to a given end point. (To find their way from start to end, the caribou use an AI path planning algorithm in which each square is weighted according to various factors such as distance to the end point and vegetation value. Finally, if the user so chooses, instead of calling up a particular year they can engage the "fast-forward" and "rewind" features in order to see a fluid cinematic representation of the changing environment.

The second mode, the hunting blind finder, is a CA which runs on top of time engine mode. For designated years, it plots markings on the 2D hashmap corresponding to the locations of simulated hunting blinds in use by simulated hunters during that year. It then generates a herd of caribou and provides them an A* path across the landscape, choosing vegetation-rich locations close to water for the start and end points of the migration path. Once the caribou have reached the end point, the blinds are scored based on various factors, and these scores help determine the placement of the blinds in the next generation. Finally, the system takes a screenshot of the hashmap representation of the environment and saves it to a bitmap file. The process repeats until the last simulation year has ended. The resulting screenshots can be used as frames for a video of the changes in the environment itself, caribou migration patterns, and hunting blind usage. Additionally, a heatmap is produced showing the percentage of time each location contained a blind that was in use versus the entire simulation time period.

The heatmap can be considered the finder's most valuable product. Its value is predicated on the fact that the more often a certain location was used, the more likely archaeologists scanning that location today are to find either paraphernalia having to do with hunting (such as speartips or spear debitage), an actual intact prehistoric hunting blind, or a related structure such as a "drive lane" (often built by hunters to funnel the caribou into a specific area. Hence, archaeologists can use it either alone or in conjunction with their intuition to determine where to send their expeditions.

Fig. 2. Our System in Action

II. Objectives and Historical Overview

A. Bottlenecks

In order to set up our experiment regarding the Alpena-Amberley Ridge, the first obvious question that must be answered is during what period of time was it actually crossable? To answer this question, we notice that the Ridge has two low-lying points that serve as bottlenecks, the first in the northwest and the second in the southeast. When the lower of these, shown on the map in Figure 1 as "α", becomes covered with water, the Ridge is not a Land Bridge, but at best two peninsulas with a strait between them. (α is about 57.5m below today's Lake Huron level, and about 118.5m above today's sea level.) When the other bottleneck, shown on the map as "β", becomes covered with water, the Ridge is at best two peninsulas with one island in the middle of them, the island separated from the peninsulas by two straits. (α is about 52.5m below current Lake Huron level, and about 123.5m above current sea level.) Now, when α becomes covered with water it is arguable that caribou could just swim the gap and continue the migration across to the other side. However, when β becomes covered with water as well, it seems clear that swimming α 's gap plus β's much larger gap is too daunting a task for any caribou herd, and hence caribou would have to give up on using the Alpena-Amberley Ridge as a migration corridor at that time, providing us with an end date for our experiment. Likewise, this reasoning also provides us with a start date for our experiment, as caribou also could not *begin* using the Alpena-Amberley Ridge until the high lake levels of the Algonquin period *receded*, uncovering the bottleneck points.

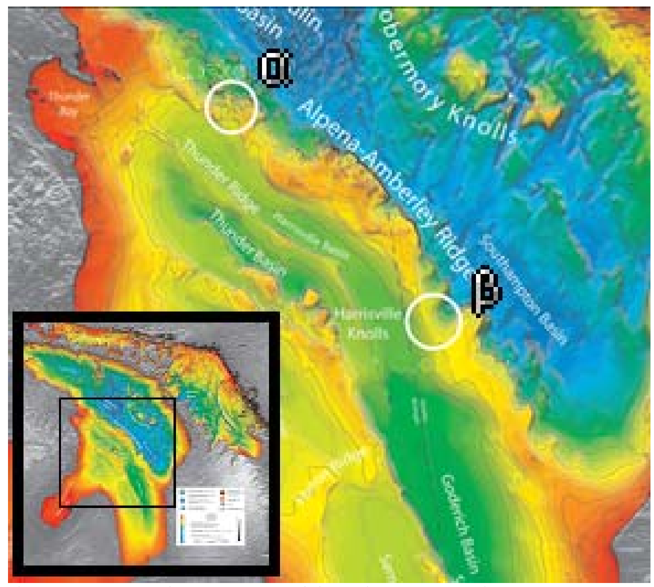

Fig. 3. NOAA Bathymetry Map of the Lake Huron Basin [2]

B. Lewis Lake Level Reconstruction

To provide ourselves with a water level time series, we use the Lewis reconstruction of prehistoric lake levels [3]. We can see from Lewis's results that in about 11800 BP, β initially became uncovered (making the Alpena-Amberley Ridge traversable), and that in 8350 BP, β as well as α became permanently submerged, never to be above lake level again. Although certain portions of the Ridge did remain above water longer, it was never to be traversable again after about 8350 BP. Given this, we acknowledge 11800 BP as the start year of our experiment and 8350 BP as the end year.

It should be noted that according to [3], there may have been three brief highstands, one at ~10600 BP, one at ~9790 BP, and the last at ~9000 BP, during which the Land Bridge was completely submerged, only to re-emerge shortly thereafter [4]. These are marked with "?" symbols on Lewis's diagram. We include these highstands in our model.

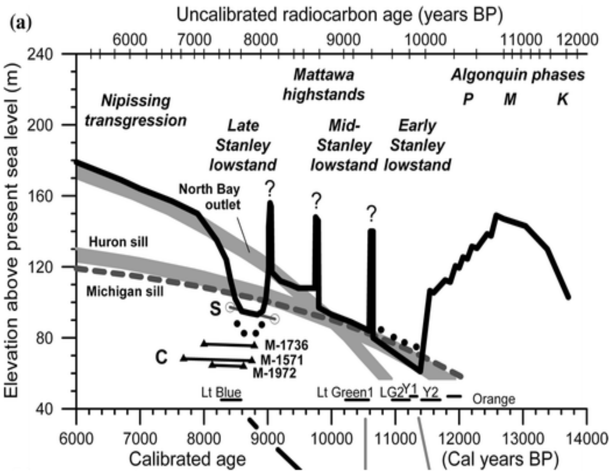

Fig. 4. Prehistoric Great Lakes Water Levels According to [3]

C. Hunting Blinds

Hunting blinds of the type Dr. O'Shea's team has found on the Alpena-Amberley Ridge are semi-permanent or permanent structures made of several large stones whose most obvious purpose was to keep animals from seeing the hunters so that the animals would unwittingly wander into spear or atlatl range where they could be killed. However, it is curious that the hunters chose stone rather than lighter materials such as wood or large mounds of dirt to build these blinds. It is quite possible that some of the hunters *did* use these materials on some occasions, but only the stone blinds have survived millennia of being underwater. Still, it is undeniable that *some* hunters chose to use large stones in lieu of lighter materials [1] [5]. A probable reason is that a blind built of wood or dirt would be washed away by a good-sized flash flood. Today, flash floods are a relatively common occurrence in certain parts of Michigan which are adjacent to the various large lakes. In addition to this, meltwater pulses at various stages of the collapse of the ice sheet may have been yet another source of flash floods back in the prehistoric era. Since the Alpena-Amberley Land Bridge has never been very high above lake level, even when lake level was at its lowest, there is no doubt that it experienced many flash floods which would destroy the temporary blinds made of dirt or wood. Permanent stone blinds would thus seem to be reserved for the most important locations and those that have proven themselves very productive over many years.

III. Cultural Algorithms and System Design

It is rare to find anything manmade as old as the prehistoric hunting blind and caribou drive lane remnants discovered by O'Shea's team. The research community is lucky to have found these few scattered remnants, yet they are not numerous enough for us to directly create a straightforward model of prehistoric hunting blind placement from their locations alone. We must thus do the next best thing, which is to use AI to simulate the hunters' human intelligence and thus their ability to decide where and when to place the blinds.

Given that humans are tribal creatures with the ability not just to individually acquire but to *share* knowledge among groups, this situation calls for a technique that reflects not just an individual but a *tribal* ability to accumulate knowledge and store it for use in future situations. In the 1970s [6], such a technique was developed by Dr. Robert G. Reynolds called *cultural algorithms.* In creating CAs, Dr. Reynolds drew an analogy between group learning, the process of Darwinian natural selection in biology, and the tendency of group knowledge acquired in the past to influence current decisions by individual members of groups [7].

A. Structure and General Algorithm

CAs contain a *population space* which is influenced by a *belief space*. Population space is defined as a set of solutions to the problem which have the ability to evolve from generation to generation. The belief space can be defined as the collected set of experiential knowledge, which has the ability to be influenced by individuals within the population space according to their varying degrees of success, and which has the ability to influence subsequent generations of individuals within the population space.

The following is a general statement of a generic CA:

1. The population space and belief space are initialized.

2. Population members are evaluated through a fitness function, and the population is ranked.

3a. The population members ranked highest are allowed to influence the belief space.

3b. In some CAs, the population members ranked *lowest* are also allowed to influence the belief space by providing *negative* information to it about their solutions.

4. The best solutions are allowed to reproduce. Operators are applied to at least some of the children which make them into mutated variants of their parents.

5. The belief space influences the children's genomes and/or their behavior in the problem space.

6. Steps 2-5 are repeated until a stop condition is reached.

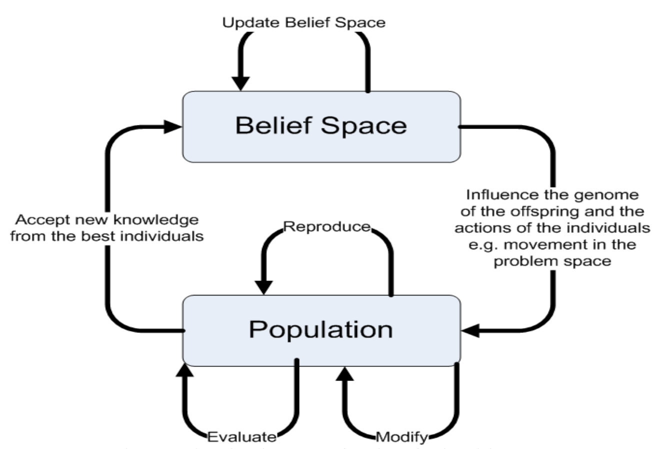

Fig. 5. Visual Schemata of Cultural Algorithms [8]

B. High-Level Design

1.) Overall Simulation: Historically speaking, it is suspected that the caribou migrated across the land bridge twice per year, once for a fall migration and once for a spring migration. However, since we intend to do 16 runs of our experiment and our experimental period is 3,450 years long, we will take into account only one "representative" migration every five years, meaning that each run of our experiment will contain 690 time intervals containing one caribou migration each and spanning five years each. We will assume also that the hunting blinds are set up once per five year period in anticipation of the caribou migration, and that they are not moved during the year once they have been set up. As described in previous sections, our system has the ability to update the terrain based on environmental change over time. To simulate the fact that the hunters cannot exactly predict how the environment is going to change, for each time interval we will first have the hunters choose locations for their blinds at the beginning of the interval (using certain weights provided by GA chromosomes, to be discussed later). Then we will use our system to update the terrain for that year, and then the caribou will migrate through. Each blind will then be scored by an objective function containing several factors which together are a fair determinant of hunt success. Each blind will then update the belief space with the values of these factors for its specific location (and, reflecting the assumption of local knowledge but non-omniscience discussed in the last subsection, the belief space will be updated with the values of these factors for each square within a 3-square Moore neighborhood of the blind as well). After this, the chromosomes for each blind will undergo various mutation operations which will produce a new batch of chromosomes (this procedure biased in favor of the chromosomes that produced the best blinds in the generation that just finished). Each chromosome will encode a "weight" for each important factor in hunting blind placement. The value of the weight for a certain factor determines how important the hunters controlling that blind consider that factor compared to other factors. Actual blind placement is determined by looking at each square in the belief space and choosing the square with the "best" values for each factor based on the weights encoded in the hunting blind's chromosome. Once the blinds are placed, yet another time interval will begin (as described before), and this process will continue until a certain terminal condition is reached. As discussed in section II B, 8350 BP is when the land bridge becomes two peninsulas separated by a strait due to rising water levels and never reconnects again. This seems the most logical choice for our terminal year for this experiment.

2.) Evaluating Hunting Blind Success and Failure: Now we need to consider precisely what factors make a blind successful or unsuccessful. One obvious factor is how close the blind is to the prey's path. Obviously a hunter throwing his weapon at prey that is closeby is much more apt to hit the prey than if he were throwing it from farther away.

Another factor is the difference between the altitude of the blind and that of the caribou. If the blind is higher than the caribou, a hunter throwing at the caribou from that blind has the advantage that gravity is working for him. In other words, a projectile thrown from that spot at the caribou below will travel faster (due to gravity) than an equivalent projectile thrown at another caribou at the same vertical level as the blind. Thus in the first case, the projectile will travel faster increasing not only the likelihood of actually hitting a moving caribou, but also the damage done to the caribou, and thus the likelihood that the caribou will be hit but still get away is decreased. Similarly, a hunter trying to hit a caribou *above* him has gravity working *against* him. Projectiles that *he* throws are more likely to miss, and what hits he makes are more likely *not* to result in killing or at least halting the caribou.
A third factor is closeness to the nearest other hunting blind. This is because hunting parties crowded too close together tend to interfere (unwittingly) with each other's success in the hunt. For example, when one hunting party makes a kill and goes to collect their kill, the herd usually reacts either by trying to move away, or sometimes even begins to panic. If the herd does panic, it is okay for the first hunting party because they have already made their kill, and all that is left is to haul it in and dress it, and then carry it back home. However, the second band, which has yet to make a kill, now must deal with a panicked herd or at least one which is wary of the spot where their fellow caribou died. This makes it much less likely that the second band will have a successful hunt. Of course on other days, the second group will make a kill first and it will be the first group which will be in the disadvantaged position. Overall, both groups will have fewer kills on average per year than if they had been spaced farther apart. Even worse, overcrowded hunting parties have an increased risk of accidentally killing each other while trying to kill caribou or other game. Obviously the farther parties are spaced apart, the less likelihood there is of a tragic accident.

Closeness to water is a factor which, although not *directly* impacting the quality of the blind location, would still impact the hunters' decision where to place the blind. Locations close to water are more likely to be flooded during the course of the year, and an underwater blind is of course useless.

IV. Implementation

A. Algorithm Development

To develop our low-level algorithm, we first constructed a preliminary version with which we ran several experiments on a small subset of Area 1 of the Land Bridge (we call this subset the "proof-of-concept area"). We used this setup to tune our constants, parameters, and algorithm components until we were receiving generally reasonable behavior from the algorithm. The main reason for using the proof-of-concept area for this task is that we were most wary of the challenge of designing a CA that could handle truly dynamic environments. Our proof-of-concept area happens to be one of the most rapidly changing, dynamic portions of all of Area 1, much more so than Area 1 taken as a whole. We figured that if our algorithm could handle the proof-of-concept area, it could handle anything else. Thus the vast majority of our algorithm was devised in anticipation of, or during, these proof-of-concept experiments, the only exception being a few small changes thought of after these experiments concluded.

B. Pseudocode

Here is the pseudocode used in our algorithm for our full experiment.

//Initialization Steps
nCaribou = 99 */*No. of caribou that cross Land Bridge each generation is 99.*/*

nHuntingBlinds = 50 */*No. of blinds (i.e., population for the GA) is 50.*/*

beliefSpace.Initialize()
*/*The belief space is an influence map with a tile corresponding to each of the regular map tiles. Each belief space tile contains four parameters, each corresponding with one of the four factors described in the high-level algorithm (closeness to caribou, height above caribou, distance from nearest other hunting blind, and closeness to water).*/*

populationSpace.Initialize(numHuntingBlinds)
*/*Initializes each hunting blind's chromosome, setting each binary bit within to a random value of 0 or 1. In this CA, each blind's chromosome consists of 16 binary bits, which are divided into 4 sets of 4 bits each. Each of these sets denotes a decimal integer which corresponds to one of the four weights that belong to this blind and help to determine its actions in response to what it believes about the environment. The four weights, in turn, relate to the distance from a given square the closest caribou approach, the height of a given square above (or below) the closest caribou, the distance from a given square to the closest other blind, and the distance from a given square to the nearest tile covered by water. (The weights themselves and the weight function process are more fully discussed in the next subsection, "Weight Function".)*/*

<u>//Main Loop</u>
*/*A logical starting point is 11800 BP, when Land Bridge is first traversable, but user can start sim at later points as well if desired.*/*

do

HBLocations = Simulation.GetHBLocs(population.genes, beliefSpace, WeightFunction)
*/*Determines the locations for the hunting blinds for this generation. If this is the first generation, locations are random. If not, locations are determined by a weight function which takes the hunting blind genes as an argument and is applied to each tile in the belief space in order to determine what the blind builders think is the most desirable spot for building the blind. (This is discussed in further depth in subsection C, "Weight Function".) In essence, each blind's new location is determined by where it* believes *the best place is according to what properties it values (which is encoded by its weights) and what all blinds collectively believe about the environment (which is encoded in the belief space).*

Also, any two hunting blinds must have at least one empty square between them. This measure was implemented to prevent the severe clumping that we saw during some of the earlier generation of our proof of concept run, where large numbers of blinds would form a solid "block" around a desirable area. Although a severe score penalty for being too close to another blind did eventually convince them not to clump directly adjacent to one another anymore, the fact still is that such close clumping would never occur in real life, which is why it is being completely disallowed for the main experiment. Note that all other aspects of the regular "closeness" penalty still apply, a blind that is only just a few (but more than one) squares away from another, although this is still allowed, still receives a hefty penalty for being too close to another blind, and the blinds eventually figure out that they must keep a reasonably healthy distance between themselves./*

Simulation.PlaceHuntingBlinds(population, HBLocations)
*/*The AI hunting teams place their hunting blinds in the locations determined by the previous function. It should be noted that our model does not include an explicit cost for constructing the blinds within the locations, however for each generation, each AI team cannot build any more than their one hunting blind built during this step. */*

Simulation.Run()

population.FitnessFunction()
*/*Computes a score for each blind based on each of the 3 important factors that determine success or failure described in the high-level design: distance to caribou, height above (or below) caribou and closeness to the nearest other hunting blind. Any underwater blind, however, automatically receives a score of negative infinity representing its uselessness despite all other factors. (Here, high scores are considered good, while low scores are considered bad.) This function is discussed in further detail in subsection D.*/*

population.SortByFitness()

beliefSpace.Update()
*/*Updates belief space tile parameters with the hunting blind score parameters (plus closeness to water) for the tiles containing the blinds, plus all tiles within a Moore neighborhood of radius 3 of the blinds (representing the hunters' ability to speculate about what might have happened had they chosen a slightly different location for their blind*/*

beliefSpace.Forget()
*/*Belief space tiles are "forgotten" (eliminated) if they haven't been updated for 10 generations (because they haven't been within a Moore neighborhood of radius 3 of a hunting blind-occupied square for that length of time). This penalizes "irrelevant" belief space squares by kicking them out of the belief space so that the corresponding map squares cannot be chosen for hunting blind placement anymore. This prevents individuals from becoming confused by very "stale" squares with obsolete beliefs and biases the hunting blind location selection process toward areas containing locations which have been successful in the recent past. */*

population.genes.Copy(exemplars)
*/*The four chromosomes yielding the belief weights that placed their respective hunting blinds in the four most advantageous locations serve as exemplars for the rest of the population. Out of the remainder of the chromosomes: 40% become copies of the 1st best blind's chromosome, 30%*

become copies of the 2nd best blind's chromosome, 20% become copies of the 3rd best blind's chromosome, and 10% become copies of the 4th best blind's chromosome./*

population.genes.Mutate()
*/*All chromosomes who are not one of the exemplars must now undergo a single point mutation. */*

population.genes.Crossover()
*/*Each chromosome is divided into two parts with a random point as the pivot. Each new chromosome is created by a random bottom part joining with a random top part. Every chromosome is subject to crossover except for the best-scoring exemplar, which goes into the next generation completely unaltered. It will be retained as exemplar (situational) knowledge for the next generation as well unless it is bested by another individual*/*

year = year + 1
generationNum = generationNum +1
map.Update() *//Updates the water level and the rest of the environment for the next year.*
until(end sim)
*/*A logical end point is 8350 BP, when Land Bridge is permanently split by relentless water rise, but user can end sim at earlier points as well if desired.*/*

//End of Pseudocode

C. Weight Function

For each new generation, each hunting blind is placed within the tile containing the highest value of the "weight function", which is applied to each "known" tile in the belief space. The value of the weight function W at belief-space tile T is calculated as follows:

$$W(T) = W_1B_1 - W_2B_2 - W_3B_3 - \text{Log}_{10}(W_4B_4) \mid (W_1, W_2, W_3, W_4 \geq 0) \quad (1)$$

In the weight function, B_1, B_2, B_3, and B_4 are T's value for the hunting blind's distance to the closest caribou approach, the blind's height above (or below) the closest caribou approach, the distance to the closest other blind, and the distance to the closest underwater point, respectively. Recalling our high-level design, we can see how the weight function is crafted so that a tile is deemed less desirable if it is far from the closest caribou, but more desirable if it has a high vantage point above the closest caribou, is far from the nearest other blind, and/or is far from water. Exactly *how much* more or less desirable is determined by W_1 through W_4, which are the weights for each of the B's, respectively. Their values are determined by the hunting blind's chromosomes. It is through these four W-values that the chromosomes determine how important the blind's builders consider each of the key factors, which together with their beliefs about the values of each key factor (represented by the B-values), ultimately determines the choice of blind construction location for the given generation.

Notice how the fourth term of the weight function, the "water fear" factor, is logarithmic. This is because if a hunting blind is already very close to water, the risk of being swamped by rising water is much greater than if the blind is quite far away. In other words, the net benefit of moving, say, 50 feet away from rising water when one is currently only 10 feet away from the water is much greater than moving 50 feet away when one is already miles away. In the former case, the benefit is crucial, in the latter, nugatory.

D. Fitness Function

The fitness score of a hunting blind H for the round is computed by the following fitness function:

$$F(H) = -C_1A_1 + C_2A_2 + C_3A_3 \text{ (if H is above water), OR}$$
$$F(H) = -\infty \text{ (if H is underwater)} \quad (2)$$

Here, A1 is the distance between H and the closest caribou approach point, A2 is H's height above (or below) the closest caribou approach point, and A3 is H's distance from the nearest hunting blind. Recalling the high-level design, each of these A's represents one of the crucial factors which determine the success or failure of a hunting party using a particular blind. C1, C2, and C3 are constants which reflect how important each of the three factors are compared to one another (1). The objective is to maximize F(H); the highest scores are considered the best whereas the lowest are considered the worst. The genes of the highest-scoring blinds are favored in the reproduction process for creating the next generation, while those of the lowest-scoring blinds are punished by being kept out of the reproduction process.

V. Experimental Framework And Results

In April 2012, we provided 400 component maps which together comprise an entire map of Area 1 of the Alpena-Amberley Ridge, the portion currently under the most intense archaeological study. Each of these component maps has 999,995 data points in it (giving a total of 399,998,000 points in all). We created these using a tool from NOAA which generates maps of the region's bathymetry (defined as the "'submarine topography,' or the depths and shapes of underwater terrain" [9]), and then interpolating in even more points. However, in our experiment, for the purpose of reducing memory usage and computation time, we are condensing all the map points into a 200 x 200 grid (40,000 points in total). In the diagram below, the larger rectangle is Area 1, the location for our experiment.

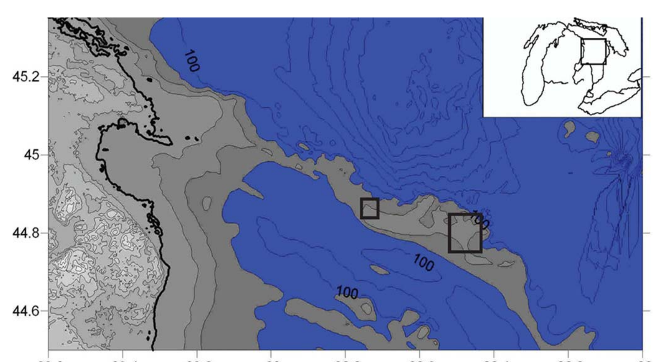

Fig. 6. Map of the Alpena-Amberley Ridge Region [1]

A. Purpose

The full experiment consists of 16 runs over all of Area 1 during the 11800 BP to 8350 BP period. Its purpose was to use our system to create a list of the locations most likely to contain a hunting blind or related artifact within Area 1. Dr. O'Shea could then use this list for his 2013 summer expedition to Lake Huron and subsequent expeditions.

B. Year Frames

Having thoroughly explained our experiment, we will now provide a few representative "year frames" from Run 1/16 to demonstrate the operation of the experiment as well as the system's ability to learn in action. In each of these frames, the white curve, the white circles, the white square, and the dark dot designate the caribou path, AI hunting blinds, highest-scoring hunting blind, and the location of the Funnel Drive Structure, respectively. The latter is a major structure containing several hunting blinds arrayed in a strategic fashion [5]. (When we informed Dr. O'Shea that we had completed a finished version of our system and were about to do evaluation runs for serious analysis, he provided us with the location of this structure. He wanted us to compare our system results specifically with this artifact, which he considers "by far the most complex hunting feature identified on the Alpena-Amberley Ridge to date." [5])

1) 11800 BP: The Alpena-Amberley Ridge first becomes a crossable corridor. Since this is the first generation, hunting blinds are placed in random non-water squares.

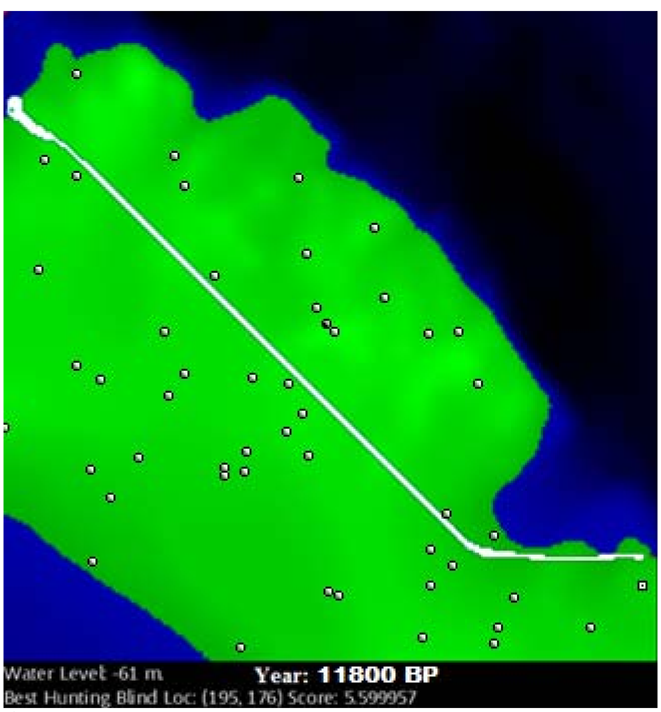

Fig. 7. 11800 BP

2) 11750 BP: Our algorithm learns very quickly. After just 10 generations (50 years), it has learned to have the hunting blinds tightly track the caribou trail. However, it has not yet learned to keep the blinds at a healthy distance from one another, and hence many of the blinds are still losing a lot of points from being too close to one another.

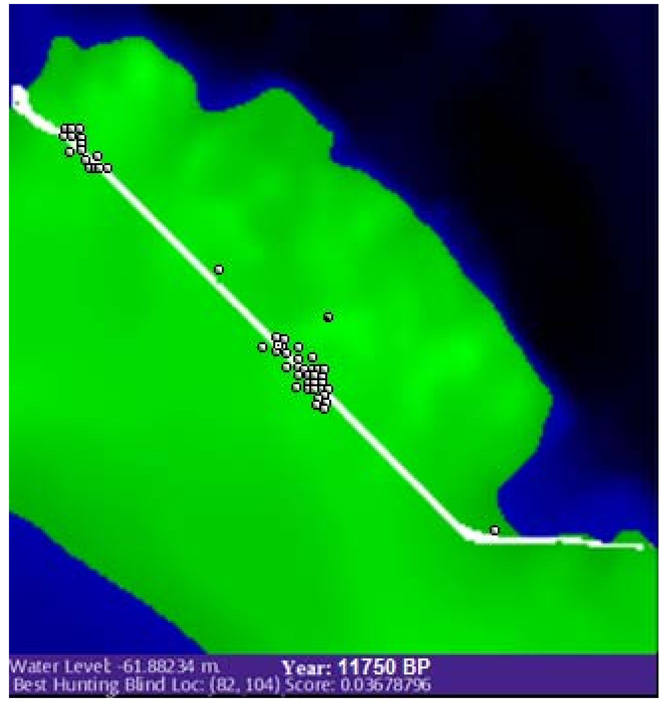

Fig. 8. 11750 BP

3.) 11230 BP: This is the middle of the Early Stanley Lowstand period, the time when the water level is lowest, the land bridge is widest, and the best spots are available for the hunters. It is during this period that the spot where O'Shea found the Funnel Drive Structure and other similar spots are most likely to have been used by the prehistoric hunters.

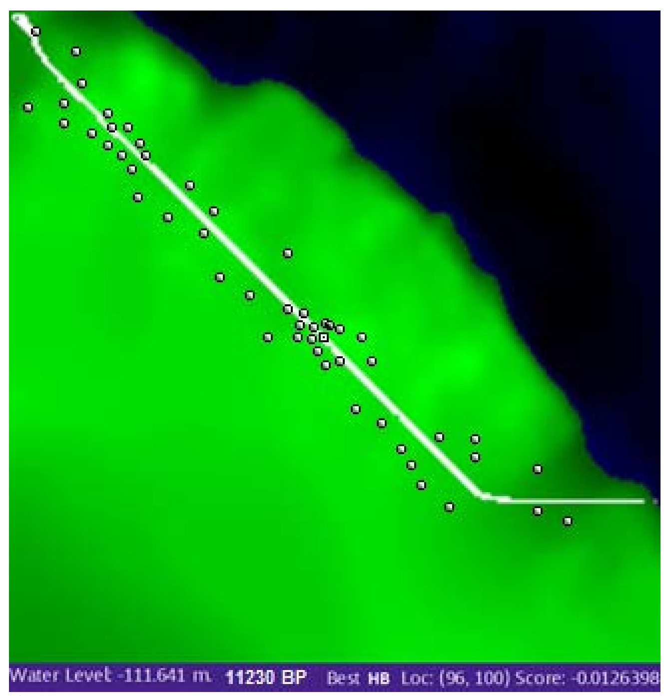

Fig. 9. 11230 BP

4) 9400 BP: This frame shows typical behavior for Mid-Late Stanley, when lake levels are very high. The caribou path is now significantly far southwest of the Funnel Drive Structure's location, so the AI blinds now no longer have incentive to go near it again. A new desirable spots now emerges on a hill overlooking the last part of the caribou path.

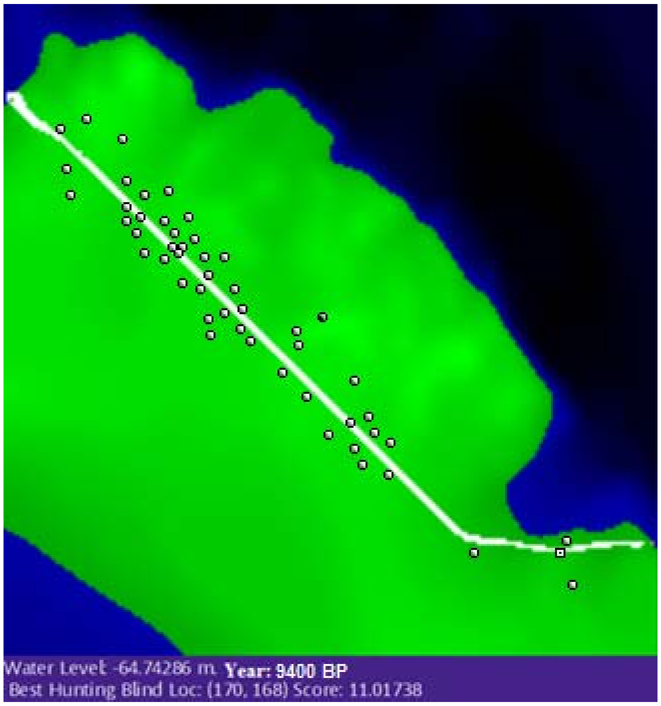

Fig. 10. 9400 BP

C. *Heatmap*

To further explore our results, we made a heatmap showing the locations most often used by our AI hunting teams.

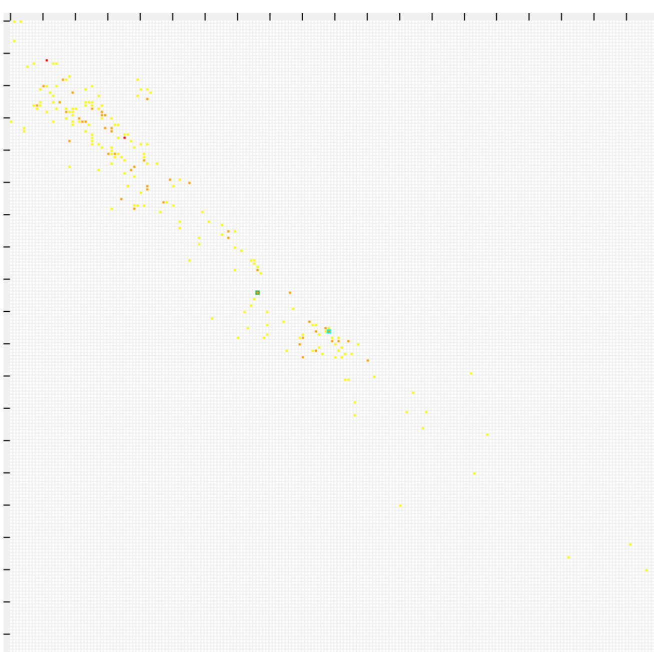

Fig. 11. Full Experiment Heatmap: Avg Hits for Each Square Over 16 Runs vs. 690 Generations (3,450 yrs)

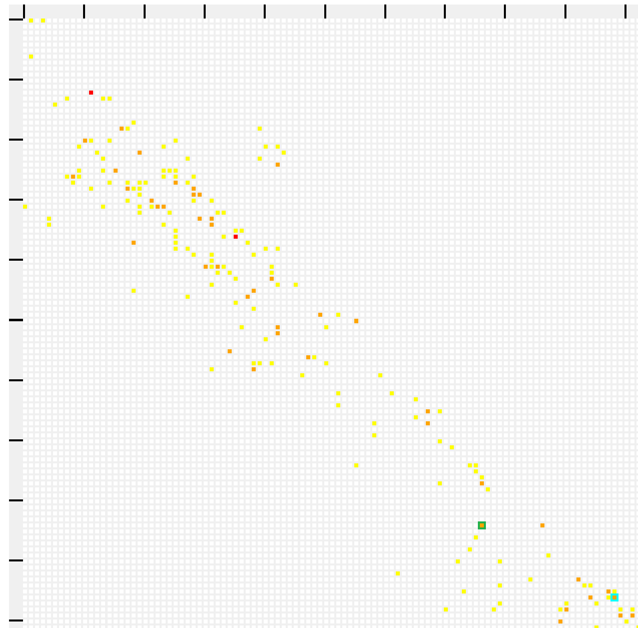

Fig. 12. Heatmap Quadrant 1

TABLE I
FULL EXPERIMENT HEATMAP KEY

Color	16-Run Average Number of Generations that this Square Contained a Hunting Blind as a Percent of 690 Generations or 3,450 Years (i.e., the Simulation Period).
Red	5%-10%
Orange	3%-5%
Yellow	2%-3%
White	< 2%
Cyan Overlay	Location where Funnel Drive Structure found. (To refrain from introducing bias, we used neither this nor any other find location as an aid in our system design. We were provided this location only *after* we finalized our model, including all parameter values used in this experiment, for purposes of system evaluation and validation.)
Green Overlay	Location where New Crossing Line discovered.

Interestingly, our heatmap analyzer found that the spot containing the Funnel Drive Structure was the 40th most frequently used location on average over the 16 simulation runs out of the 40,000 locations contained in the grid representing Area 1.

Dr. O'Shea was very intrigued by these results. He also mentioned that he was especially interested in another one of our locations which he had not searched yet. During his next expedition season (in the summer of 2013), he made a stop to that location and found a very large line structure which he named the New Crossing Line. He believes there is a good chance that this was a caribou "drive lane" meant to funnel caribou toward where the hunters lay in wait for them.

D. *Learning Curve*

For convenience, we now provide a learning curve for our algorithm.

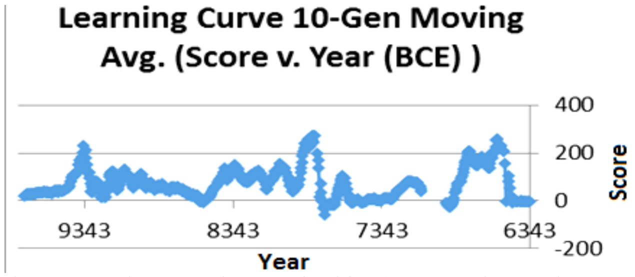

Fig. 13. Learning Curve for Our Algorithm (10-Generation Moving Avg. of Highest-Scoring Hunting Blind Score vs. Year)

The CA learning curve seen here is unlike most other CA learning curves, however there are important reasons for that, the biggest being that our objectives are not static. Caribou paths, and most importantly water levels, are subject to sudden and unpredictable change. What had been a great hunting spot for a few or even many generations may not be so good, or may be completely unavailable, the next generation. Also, the four major catastrophic water rises which befall the land bridge are major hampers on learning because they create significant periods in which the caribou do not even attempt to cross the land bridge, creating a major disruption for the hunters. Nevertheless, we can see that the algorithm is indeed learning. Notice how the 10-generation moving average reaches its overall peak during Mid Stanley, even though the water level is lower (and hence more hunting spots are available) during Early Stanley. Notice how even the Late Stanley peak for the 10-generation moving average is higher than for the Early Stanley period, even though the water level is significantly higher in Late Stanley than Early Stanley. It is only during Mid-Late Stanley, when the water level is extremely high and there are many fewer good hunting spots available than in the other periods, that the peak fails to exceed that of the Early Stanley period.

VI. Intelligent Caribou

Despite the success of our algorithm's ability to predict hunting blind locations, Dr. O'Shea was able to identify some locations bearing artifacts that were not among the most highly recommended positions by this system. He suggested that accounting for multiple paths to the south and west of the original experiments could solve this problem.

While work on the hunting blind prediction algorithm was taking place, another researcher with the Land Bridge Group, Jin Jin, was developing a CA based path planning algorithm for the caribou herds [5].

We took O'Shea's advice and adjusted the simulation to allow for multiple simultaneous paths. Also, in order to make our model's caribou paths more accurate, we also replaced our AI caribou (which contain the A* path-planning algorithm) with Jin's AI caribou (which contain his CA path-planning algorithm) [5]. Also, because Dr. O'Shea's expedition team was especially interested in artifacts the period of 9398 BCE to 9237 BCE, we also decided to perform our new batch of experiments using that year range, Because the interval is much smaller we made our time step to be 1 year instead of 5. These experiments required more processing power because it was running two CA's and, for the same reasons mentioned in section V, the data had to be scaled down even further so that the map contained 80 X 80 data points, 6400 total.

We ran preliminary experiments, and immediately saw interesting results. The three caribou herds achieved a confluence in their migration pattern as seen in figure 12. We also saw a large concentration of predicted locations in this confluence area. These results were presented to Dr. O'Shea who was then able to discover artifacts and debitage in this area of confluence.

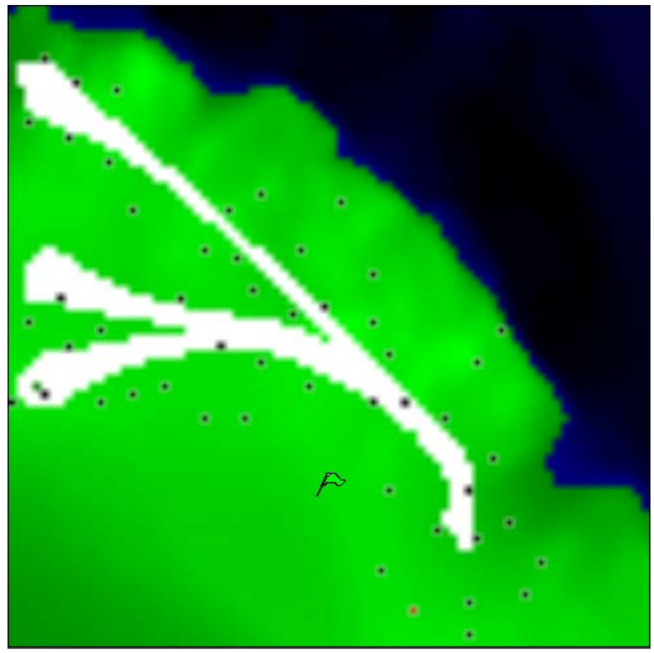

Fig. 14. Hunting Blind Simulation using CA Path Planning for Caribou Herd Migration.

For convenience, we also provide a heatmap displaying the overall results of our experiment.

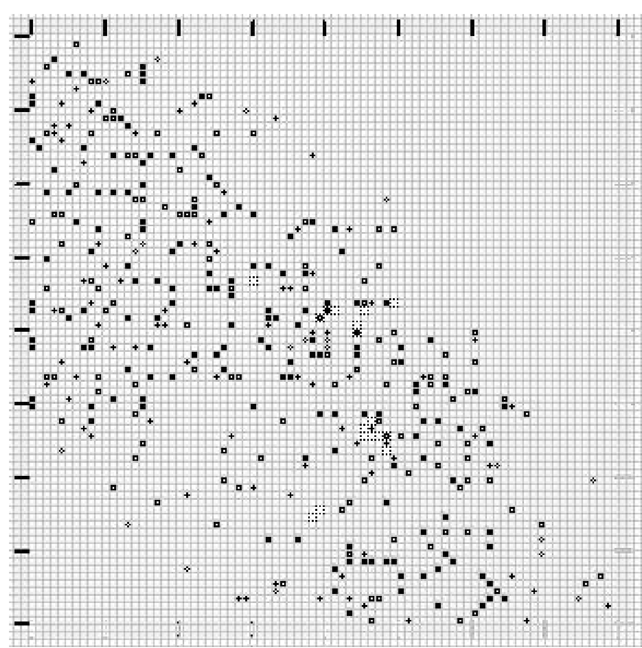

Fig. 15. Heatmap for Hunting Blind Simulation using CA Path Planning for Caribou Herd Migration.

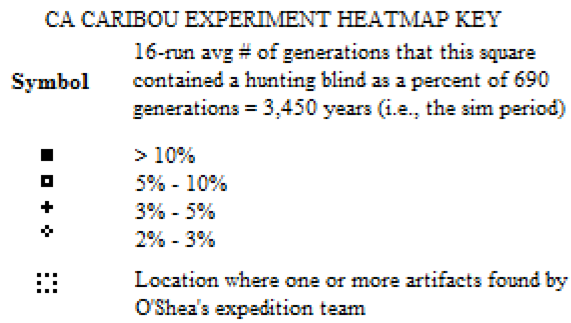

Fig. 16. Heatmap Key for Hunting Blind Simulation using CA Path Planning for Caribou Herd Migration.

Our heatmap in Figure 15 seems somewhat sparse, likely because in the region before the herd paths converged, the hunting blinds were dividing themselves amongst the three separate herds in order to hunt all of them (as was proper). Nonetheless, there is a key overall pattern here. In order to better display it, we have created a 40x40 heatmap by consolidating the locations in Figure 15.

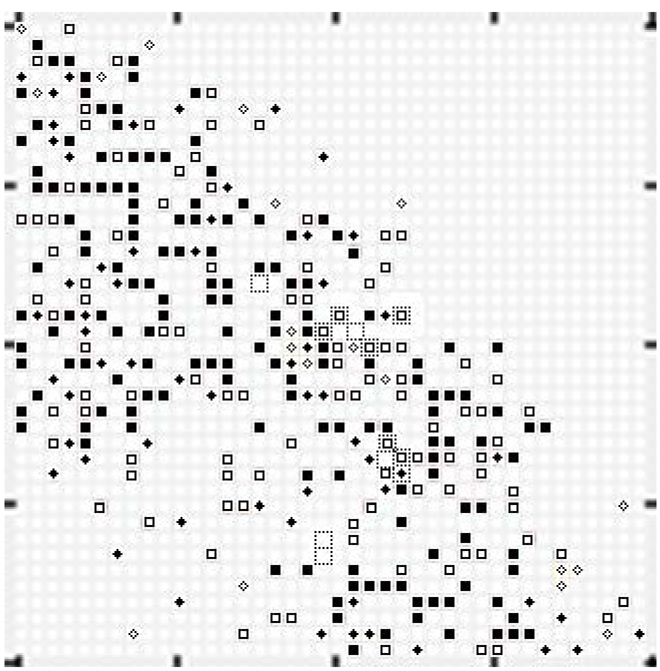

Fig. 17. Heatmap for Hunting Blind Simulation using CA Path Planning for Caribou Herd Migration.

Here one can easily see the distinctive arc pattern created by blinds trying to track the caribou herds as they enter from various points in the west, converge toward a narrow confluence region, and then exit Area 1 by heading south. Looking at these results, the narrowest part of the confluence point would seem the next logical place for archaeologists to explore.

VII. CONCLUSION

Again, our core method is by no means limited to just prehistoric hunting blinds and other objects directly related to them such as caribou drive lanes and spear debitage. Given a set of rules about where any type of artifact can generally be found with respect to various conditions and features within its environment, our CA system can incorporate those rules, our time engine can simulate the changing environment during the relevant period, and heatmaps of object locations over the time period according to the simulation can be produced, just as was the case in the hunting blind experiment featured in this paper. The next step in our research will be to run our system over the entire Alpena-Amberley Land Bridge in order to produce a heatmaps encompassing the entire land bridge that archaeologists can use to decide which areas they should spend their time in searching for hunting blinds and other related artifacts, and which to ignore.

REFERENCES

[1] J. M. O'Shea and G. A. Meadows, "Evidence for Early Hunters Beneath the Great Lakes," *National Acaemy of Sciences Proceedings,* vol. 106, pp. 10120-10123, 2009.

[2] National Oceanic and Atmospheric Administration, "Great Lakes Bathymetry," [Online]. Available: http://www.ngdc.noaa.gov/mgg/greatlakes/greatlakes.html. [Accessed 2012].

[3] C. F. M. Lewis, C. W. Heil, J. B. Hubeny, J. W. King, T. C. Moore and D. K. Rea, "The Stanley Unconformity in Lake Huron Basin: Evidence for a Climate-Driven Closed Lowstand About 7900 14CBP, With Similar Implications for the Chippewa Lowstand in Lake Michigan Basin," *Journal of Paleolimnology,* vol. 37, no. 3, pp. 435-452, 2007.

[4] C. F. M. Lewis, S. M. Blasco and P. L. Gareau, "Glacial Isostatic Adjustment of the Laurentian Great Lakes Basin: Using the Empirical Record of Strandline Deformation for Reconstruction of Early Holocene Paleo-Lakes and Discovery of a Hydrologically Closed Phase," *Géographie Physique et Quaternaire,* vol. 59, no. 2, pp. 187-210, 2005.

[5] J. O'Shea, A. Lemke and R. G. Reynolds, ""Nobody Knows the Way of the Caribou": Rangefier Hunting at 45° North Latitude," *Quatrenary International,* vol. 297, pp. 36-44, 2013.

[6] R. G. Reynolds, *An Adaptive Computer Model of the Evolution of Agriculture for Hunter-Gatherers in the Valley of Oaxaca, Mexico,* Ann Arbor: PhD Dissertation, University of Michigan, 1979.

[7] R. G. Reynolds, "An Introduction to Cultural Algorithms," in *Proceedings of the 3rd Annual Conference on Evolutionary Programming*, 1994.

[8] J. Jin, *"Path Planning in Reality Games Using Cultural Algorithms: The Land Bridge Example.",* 2011.

[9] National Oceanic and Atmospheric Administration, "National Ocean Service," [Online]. Available: http://oceanservice.noaa.gov/facts/bathymetry.html. [Accessed 2014].