The Deep Dive System: Reimagining Global Climate Change with Artificial Intelligence and Virtual Reality

Sarah Saad[1], Thomas Palazzolo[2], , Chencheng Zhang[3], Robert G. Reynolds[4,6], Ashley Lemke[5] , John O'Shea[6] and Calien O'Shea[7]

Wayne State University[1,2,3,4], University of Texas[5], University of Michigan Ann Arbor[6], North Dakota State University[7]
fb5308@wayne.edu[1], dz3786@wayne.edu[2], gl4705@wayne.edu[3], robert.reynolds@wayne.edu[4], ashley.lemke@uta.edu[5], joshea@umich.edu[6], cailen.oshea@ndsu.edu[7]

Abstract

The Deep Dive system was designed originally to predict ancient paleolithic site locations using AI and Virtual Reality. The focus of the project was a Land Bridge that connects what is now Michigan to Canada. During the recent Ice Age it was above water from 10,000 to 8000 B.P. During that time, it was a migration route for caribou. The assumption is that where caribou go, hunters will follow. AI path planning approaches were used to predict caribou migration pathways over time on the Land Bridge. This work investigates the impact that herd size has on the migration routes taken. Three models of herd behavior are used by an evolutionary algorithm, Cultural Algorithm, to produce optimal values for each of the herd models for different sized herds. The question of interest is whether there was a "tipping point" for one or more of the models that lead to a bifurcation of the optimal path over time? The results suggest what models are most appropriate to answer this question but the reasons why the bifurcation took place.

Keywords—Tipping points, Machine Learning, Evolutionary Algorithms, Cultural Algorithms, Virtual Reality.

I. Introduction

While it is possible to observe tipping points for environmental processes over the short term, long term evolutionary change is much more difficult to assess. The Deep Dive Land Bridge simulation system was initially developed to aid underwater archaeologists in the discovery of ancient ice age prehistoric occupational sites now underwater in Lake Huron, one of the Great Lakes in the United States (Drikx et al, 2013)(Amato, 2017)(Reynolds et al, 2011). It utilized Artificial Intelligence and Virtual Reality to recreate the archaic semi-artic landscape and has facilitated the discovery of several ancient underwater sites (Fogarty et al, 2015).

The Land Bridge was above Lake Huron water level for about 2,000 years from 10,000 B.P. since water levels fell because of glacial formation during the Ice Age. It was 8 miles wide and 80 miles long and connected what is now Alpena in Michigan to Amberley in Ontario Canada. The region was semi-artic with many plants and animals that are no longer present. For example, woolly mammoth skeletons have been recovered in Michigan. As the global climate began to warm, water was released from the nearby glaciers and the lake level rose. By 7600 B.P. the Land Bridge was no longer above water.

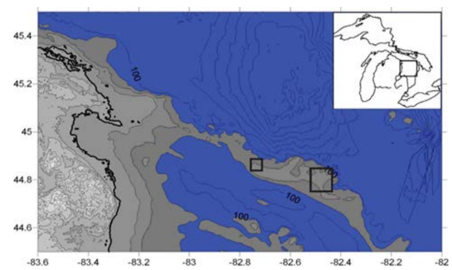

Figure 1. The Land Bridge extends from Alpena, Michigan to Amberley Ontario on the east.

Figure 1 shows the location of the Land Bridge relative to the current state of Michigan on the west and Ontario Canada on the east. The two cells on the bridge represent areas that are the focus of original archaeological exploration. They were selected due to their location relative to the widest part of the Land Bridge.

The key component of the Virtual Reality system are the caribou herds and their migration patterns through various biomes. Ancient hunter behaviors would have related to their ability to predict herd movements and exploit those predictions. Figure 2 shows a caribou herd moving across the Land Bridge VR landscape.

Figure. 2 A caribou herd moving across the Land Bridge.

While the Land Bridge was above water it could serve as a migration pathway for thousands of caribou in the spring and fall. At that time caribou were a major food source for Paleoindian hunters. As lake levels rose there was a point when geologic, hydrologic, and environmental factors worked together to produce a "tipping point" events. One such event was when the land bridge was no longer an isthmus connecting two larger land masses (Lewis, 2016).

A second more local "tipping point" is a behavioral one in which caribou herds created different patterns of crossing behavior. In order to address these three A* path planning algorithms were developed based upon different models of herd behavior. For small herds an optimal path favored the southern edge of the bridge. However, ancient structures have been found on the northern side of the bridge. The question of interest here is whether there a "tipping point" for each optimization algorithm at which the optimal pathways branches both to the north and south based upon herd size and the environmental parameters?

The organization of the paper is as follows. In section 2 several models of herd movement observed in the real world are described: Single Path A*, A*mbush, and Dendriform A*. Each of them will be used to generate optimal migration pathways across the Land Bridge for different herd size categories. Next in section 3 simulations are performed using basic calorie content information to identify three categories of herd size based upon the amount of STRESS placed on the herd as it moves across the landscape. The goal is to identify the number of animals in herds with high survival rates, average survival rates, and low survival rates respectively. Then in section 4 Cultural Algorithms, an evolutionary hyper-heuristic, is briefly described and the employed to generate an optimum pathway for each herd size algorithm combination. Section 5 provides the results about the tipping points observed and the suggested reasons behind them. Section 6 gives the conclusion.

II. The Multi-Agent Planning Framework for the Deep Dive Simulation Component

Figure 3 gives an overview of the overall Deep Dive system. It has three basic components: The Pathfinder MAP Simulation system; the Graphical User Interface (GUI) for the simulation system; and The Virtual Reality system.

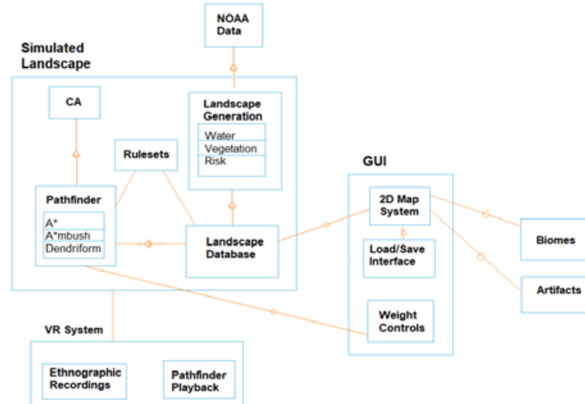

Figure. 3 The Overall Organization of the Deep Dive system.

The topographic data acquired from the National Oceanographic and Atmospheric database (*NOAA*) of the area was fed into the AI pipeline to *Generate* AI content via the Landscape. This created content includes the water level of various cells of the landscape to identify which areas of the Land Bridge were above the current water level or not for a given year between 10,000 and 8,000 B.P. For any given year height map data for those portions of the landscape was calculated along with derived slope. Hydrological information including the location of ponds, swamps, and rivers that are present in the location were then calculated. Given the location, water content, slope and sun angle the AI pipeline can predict the cells potential vegetation at each location on the Land Bridge. This information is stored in the *Landscape Database* for use by the *Pathfinder* system.

The basis for the simulation system is *Pathfinder*, a Cooperative Multi-Agent Planner (CMAPP). There are several deterministic general purpose MAP solvers available (McLeod et al, 2017).. They include MAPR (MAP Planning by Reuse), CMAP (Cooperative MAP), mu-SATPLAN (Satisfiability based planning), among others. The different CMAP solvers are be classified by the mechanisms that they employ to address the planning process. The main features that can be used to characterize cooperative MAP solvers are:

1) Agent Distribution: The MAP process here involves multiple agents who are involved in the planning process either as active participants or as target for the planning process.
2) Computational Process: Whether the computational process is performed using a

centralized monolithic processor or distributed among several processing units.

3) Plan Synthesis: This involves how and when the coordination activity is applied among agents. Coordination activities represent how information is distributed among agents and how their actions are combined together.
4) Communication Mechanism: How agents communicate with each other.
5) Heuristic Search: MAPs that use local heuristics to allow individual agents to assess their estimate progress towards their individual goals. Those with global heuristics calculate them for all the agents.
6) Privacy Preservation: Multi-Agent problem solvers can be distinguished in terms of their use of various privacy algorithms.

The CMAP, *Pathfinder*, used here was developed especially for the computational needs of this project. It is a monolithic, hierarchical, and Multiagent Planner based upon the A* Algorithm with the caribou agents as the target of the planning process. The planner uses a global heuristic to generate a single A* optimal path. This optimal path is used as basis for the A*mbush algorithm that decomposes the original path into waves of agents. Each wave consumes a certain amount of resources leaving the remainder for the next wave. The number of waves is given as a parameter. Then the results are given to Dendriform A* which decomposes the waves into smaller subgroups based upon environmental parameters. The result is to generate a set of two-dimensional waypoints that support the optimal path across the Land Bridge.

The three algorithms comprising the Pathfinder approach are now briefly described:

Single Path A*: A* is a popular search-based pathfinding algorithm that's an adaptation of Dijkstra's Algorithm. The difference being an additional heuristic allowing it to attribute cost to actual and estimated distance from the goal. Since the algorithm calculated point by point it allows the caribou agents to traverse the landscape while focusing on effort, risk, and nutrition.

Algorithm 1: A* Pseudocode

```
Add pathStart to openNodes
    Initialize pathStart scores
    While (openNodes count greater than 0)
    {
        currentNode = openNodes [0]
    If (currentNode is goalNode)
    { assemblePath() and return true}
        Remove currentNode from openNodes
        Find currentNode's neighboringNodes.
        ForEach(neighboringNode) calculate f and g score
        If (neighborNode is not in openNodes) add to
openNodes.
        Else {adjust neighborNode's position in the
openList based on total score}
    }
```

A*mbush: Another migration algorithm integrated into the system is A*mbush. A*mbush incorporates A* at its root. It uses the algorithm of A* but does so in separate waves instead of a single path. The number of waves are entered as a parameter, then the total herd size of Caribou is divided amongst the waves. The waves are then sent one after another with the nest wave entering the landscape as the last one completes its journey. Each wave consumes a certain proportion of available calories, leaving the remainder for the waves that follow.

Algorithm 2: A*mbush Pseudocode:

```
for (generations=0; generations < A*mbushGenerations;
generations++)
for (waypoint = waypoints-2; waypoint > 0; waypoint--)
{foundPath = AStar(waypoint, waypoint+1)
foreach(node in foundPath)
{Insert node in resultPath(generations) at index 0.}
} resultingHerd += calculateMigrationScore()
devourVegetation(foundPath)
}
```

Dendriform: Dendriform is the final algorithm used in the path planner portion on the Deep Dive system. It incorporates A* at its root but also allows for branching during the exploration of the landscape. This means as the line of Caribou is traversing, they can divide on the spot allowing some of the herd to continue their path while the rest look for a separate path. The subherds can coalesce in order to reduce risk when calorie intakes are satisfied.

Algorithm 3: Dendriform Pseudocode

```
Calculate optimal A* path
Add starting point and ending point to node list.
While node list has more than two nodes {
checkForNewDivergencePoints
select last two nodes in node list and A*mbush Devour
path section.
Remove last node in node list.
If last node in node list is not starting point:
    Calculate optimal A* path to ending point.
}
```

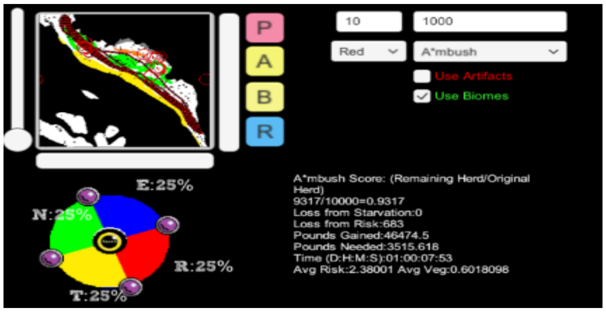

Figure. 4: The Current Deep Dive GUI (Left side of screen). The upper left gives the path produced by an algotihm. Bottom left gives the 4 basic parameters for the optimization.

The Simulation system then communicates with the simulation GUI in two basic ways. First, the user interface displays a series of tabs through which the user may navigate to a given data set or select an experiment to run as shown in Figure 4. Maps can be viewed in a variety of data styles, such as biome data, topographical data, archaeological points of interest, ruleset hotspots, and so on. Pictured above, the user has selected to run six iterations of the A*mbush pathfinder, each wave being made up of 1000 caribou. The weight priority wheel on the bottom left allows the user to manually set the weights for Effort, Risk, Nutrition, and Time in the performance function. The priority weights control what will be important to the caribou in the current run. The green segment denoted by a "N" is the nutrition this will have caribou prioritize situations which will lead to an increase in calorie or food intake. The blue segment is effort ("E"). An increase in this priority will cause the caribou to avoid scenarios that lead to excess calories being spent for example going up a steep incline. The red segment is risk ("R") which influences caribou to avoid scenarios that lead to a higher percentage of deaths. The last weight denoted by yellow is time ("T"). This parameter prioritizes the amount of time it would take to cross the entirety of the portion of the land bridge simulated.

In the next section, the three algorithms will be used to generate survival scores for a range of herd sizes from 50 to 250,000. Three representative herd sizes will be extracted from the resultant distribution to be used for the optimization phase.

III. The Prototype Herd Movement Model

In this section the goal will be to observe the survival scores for each of the three algorithms over a range of herd sizes. Representative herd sizes will be extracted from the resultant distribution for use in the optimization phase.

In the following, the model assumes a Fall migration pathway over the Land Bridge under the control of the four basic parameter components: Risk, Nutrition, Effort, and Time. To better visualize the tradeoffs between the components STRESS charts were produced as a result of simulating the optimal path constructed by each algorithm for herds ranging in size from 50 to 300,000 across the Land Bridge. Herd size is plotted on the x-axis and the herd survival as a percentage is plotted on the y-axis. Figure 7 gives the chart for the Dendriform algorithm.

Figure 5 combines the STRESS curves for all three algorithms under the "all things being equal" assumption for the weights. In that Figure A*mbush (1 wave) is the same as A*. Notice that there are three basic phases. In the first phase, survival rates increase as herd size increases from 50 to around 8000. This reflects the principle of safety in numbers. The scenario that prefers risk over all other factors dominates in that phase. The next phase between 10,000 and 20,000 represents a plateau where the impact of adding new members is offset by an increase in members who are lost due to starvation. In the final phase the nutritional concerns start to dominate with herds above 25,000. In that phase, the herd models that are able to distribute individuals over the landscape are best able to ameliorate the observed reduction in survivability.

It can be seen that the more the herd is broken into waves the lower the slope for risk reduction. A* has the highest rate of reduction in risk since all of the individuals in a herd are used at once. Dendriform is a close second since all of the individuals start together although they separated into groups to achieve higher caloric intake.. The plateau phase is the shortest for A* followed by Dendriform. For A*mbush the more waves there are, the longer the equilibrium phase and the shorter the nutrition dominant phase.

Figure. 7: The GUI screen used to display and modify heuristic constraints on generated paths

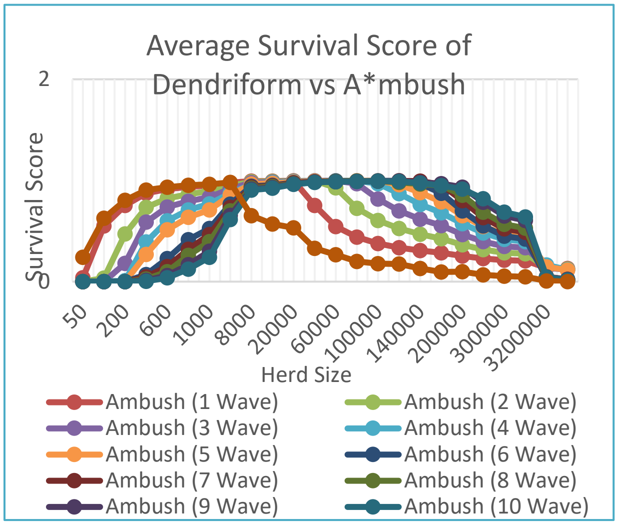

Figure. 5: Survival rates equalizing all weights using the A*, A*mbush and A*Dendriform algorithms across various herd size. A*mbush (1 Wave) is the same as A*.

From Figure 5 three herd categories can be identified. Herds of size 8000 are representative of high survival herds where risk is a dominant parameter due to their relatively small size. In other words, the larger the herd the less accesible they are to predators such as wolves. So increased herd size is a principle factor affecting survival for small herds. Herds of 15,000 are selected to represent scensarios where reduced predation is counterbalanced by increased need for food resources. A herd size of 25,000 will represent herds where caloric content is the dominant factor in survivabilty. These will be the three representative herd sizes to be used in the subsequent path planning optimizations.

IV. Using Cultural Algorithms to Generate Optimal Caribou Migration Paths.

Here a machine learning algorithm, the Cultural Algorithm (*CA*) is employed to produce a set of weights that optimize group survivability. The Cultural algorithm is a socially motivated algorithm developed by Reynolds (McLeod et al, 2017), (Culturally Responsive-Sustaining Education Framework). It's a means to solve problems in a complex system like the ones posed to the Deep Dive's path planner. It is described graphically in Figure 6. The CA is composed of a belief space and population space. Here the population is a set of experiments that employ different values weights for environmental parameters. The knowledge sources are housed in the belief space and represent the acquired knowledge of the population.

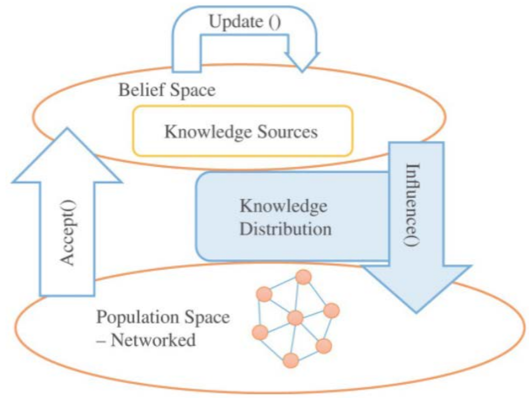

Figure. 6: Cultural Algorithm Representation.

The knowledge sources are learned from the experience of the population of problem solvers. The basic categories are exemplar individuals, current range of acceptable variable values, rule constraints, and prior performance history.

Individual agents in the population are then influenced by the belief space knowledge and are linked together in a social network. In each generation every agent has its current strategy, but it is able to see those of its immediate neighbors. A weighted majority vote is taken that includes the individual and it neighbors. The winning strategy is then used for that individual. The agents resultant decision is evaluated by their relative fitness. Top performers are accepted into the belief space and used to update the knowledge there. Here the knowledge sources determine the optimal weights priorities of risk, nutrition, time, and effort can have on the system.

In addition, the user can place constraints on the generated pathways in two ways. First, they can require the generated path to be constrained to pass through a set of manually set waypoints. Those points can be set by clicking on the map in the upper left-hand corner of the screen. Also, users can select rules that constrain the regions within which the paths can be placed. The Rule Selection screen is shown in Figure 8 below

The Cultural Algorithm was applied to produce the optimum component weights for the three representative herd sizes determined earlier. One herd size was selected from each of the three phases in the model. The biome information used in these tests was the Sonnenberg Version 3 biome map, created from a collection of regional polygons.

Figure 8 gives the color-coded key for each of the biomes. The presence of water bodies are indicated in black. The white color is the basic tundra, a low-hazard traversable biome. Green represents marshlands, abundant in nutrition but high in risk and requiring more time to move through. Yellow is the sandy beach biome, with relatively low risk but low nutrition and a time penalty for moving in soft sand. Orange represents the rocky biome, with lower vegetation and light risk and time penalties due to the uneven, rocky surface of the landscape. Gray represents the northern cliff biome, which is immediately adjacent to a vertical drop which would likely be fatal to any caribou that fell from it. The dark red areas represent eskers. Eskers are long ridges of gravel and other sediment with a typical winding course that are deposited by meltwater from a retreating glacier or ice sheet. The biome color code key is also given in Figure 8

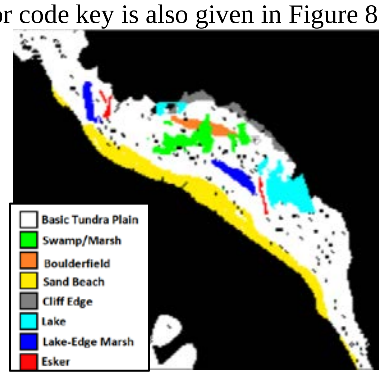

Figure. 8: Current version of the biome data comprised of polygonal regions which includes the lake biomes, the lake-edge marsh biomes, and the esker obstacles among others. Color coded key to the biome map is also included.

The Cultural Algorithm was then applied to produce an optimal configuration of weights for each of the three representative herd sizes when run with A*mbush and Dendriform. A* was omitted since it was a special case of A*mbush with just one wave. The algorithms were applied to the Land Bridge that was above water from 55 through 60 m below sea level at one meter increments. One area common to all these landscapes is the focus of the results presented here.

V. Optimization Results

The Cultural Algorithm was used to produce optimal paths and parameter weights for three algorithms across all 7 lake levels, 55 through 60 m below sea level. across the entire Land Bridge. There was one area that was in common with all of the lake levels examined defined by the 40m contour from the lake bottom. That area is shown in Figure 9 along with the location of known hunting structures.

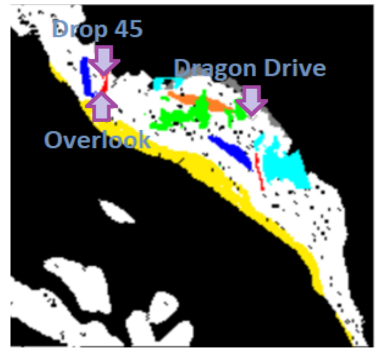

Figure 9: Locations of known hunting structures.

Note that two of the structures, Drop 45 and Overlook, are positioned at either end of the esker ridge. The esker is an obstacle that caribou need to go around. The third structure, Dragon Drive, is located on the northern side of the bridge and may have been used as a location to which caribou were driven.

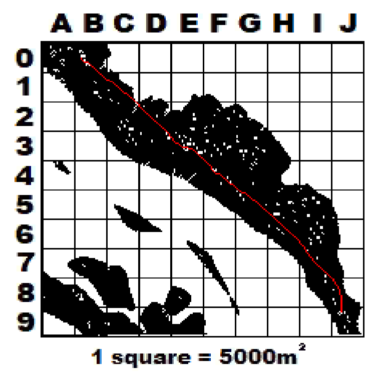

Figure 10: Dendriform A* path for a herd size of 8000.

	A*mbush 8	**A*mbush 15**	**A*mbush 25**	**Dendri 8**	**Dendri 15**	**Dendri 25**	**AVG**	**STD DEV**
Effort	19	20	33	37	33	8	25	10.18168
Risk	12	16	11	24	6	3	12	6.806859
Nutrition	61	46	46	32	46	79	51.66667	14.81741
Time	8	17	10	8	15	10	11.33333	3.448027

Table 1: Optimal parameter values for each of the algorithm herd size combinations,

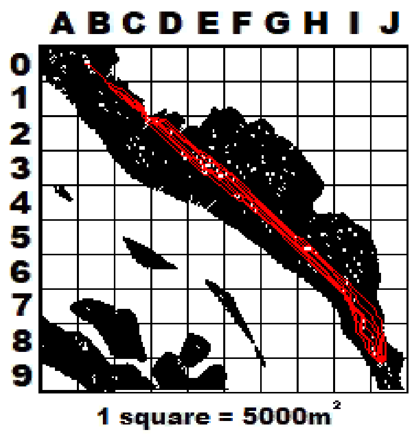

Figure 11: Dendriform A* for a herd size of 15000.

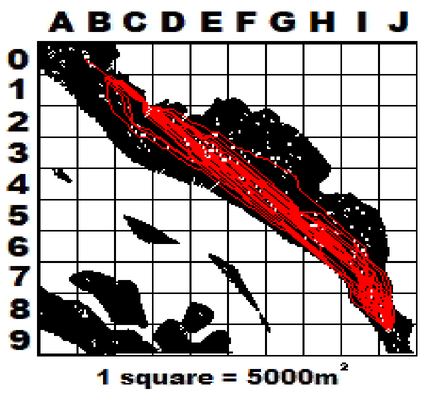

Figure 12: Dendriform A* for a herd size of 25000.

The key question here is whether there is a tipping point in terms of herd size when there are enough caribou to traverse both the north and south areas? If so, is there a different tipping point for each herd model and what are the optimal values for the four components?

Figures 10 through 12 present the paths generated by Dendriform A* for each of the three categories. The row delimiters extend from the west border on the left to the east border on the right. Likewise the column delimiters extend from the northern border at the top to the southern border at the bottom. The paths for both the 8000 and 15000 herd sizes tend towards the southern side of the bridge. Only high frequency paths are represented in these diagrams. However, for a herd size of 25000 there is a substantial northern component. There appears to be a tipping point between 15000 and 25000 that accounts for that shift but the reasons for the shift are different in each case.

To identify the reasons for the shifts one can look at the optimized parameters for the algorithms for each of the three population sizes. The optimized parameters are given in Table 1. For A*mbush the optimum values increase for effort, risk, and time spent whereas that for nutrition decreases substantially. The traditional southern pathways are becoming more challenging for the herd. For a herd size of 25000 the time spent, and risk factors were reduced while the emphasis on nutrition stayed the same. Only effort increased its share since each wave needed to spread out more because there is less food left for them by preceding waves. By distributing the herd in a more northerly direction pressure on overall movement appears to have been reduced.

While A*mbush herds move out in waves Dendriform A* begins as a single herd and breaks into sub herds to generate more nutrition. Nutrition is a key factor as herd size increases. Expansion to the north reduces risk since sub herds need not break up so often and can be larger in size. The emphasis on effort and time are also reduced because less congestion improves speed of movement and less effort is needed to coordinate group fission and fusion.

It appears that the tipping point in terms of movement to the northern side of the bridge could occur with either herd model but for different reasons. For the wave based models it is due to the reduction of caloric intake as a result of many waves moving over the landscape. On the other hand, for the sub herd based model it allows more exploration of the environment but with larger sub herd sizes so as to reduce risk and expedite movement.

In addition, what the results do suggest is that the northern most locations were most viable for larger herd sizes, those most likely in the fall. Thus, this behavioral tipping point for the caribou may also have produced a corresponding tipping point for the paleolithic hunters. It is possible that these northern locations were constrcuted after the structures found on the southern side. However, once the glaciers began to melt and the lake levels rose again the size of the herds would be correspondingly reduced. At the same time the southern locations would be progressively under more water and the remaining herds of smaller size would be attracted to the northern region. Thus, the northern most sites developed later but lasted longer than the southern ones until the lake levels reached a tipping point and caribou crossing was no longer viable.

Figure 13 overlays the frequency of traversal on the virtual land bridge generated by all of the runs overall herd sizes and algorithms. The view is from cell D2 in Figure 11 towards the southwest. Red cells were traversed by 100% of all runs. Yellow cells were traversed approximately 75% of the time. Light green and blue cells were traversed slightly above and below 50% respectively. Dark blue and black cells are close to 0%. In summary, the use of artificial intelligence to generate and test hypotheses concerning behavioral tipping points has provided the opportunity for a deeper understanding of the impact that climate change has on the environment on a behavioral level.

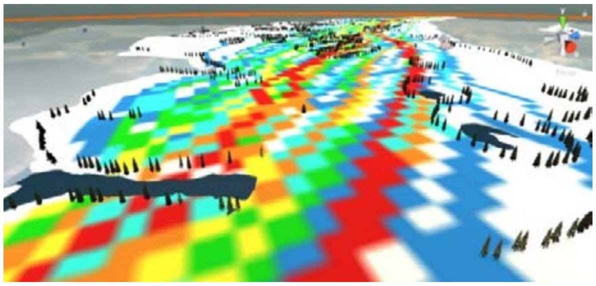

Figure. 13 A view of migration frequency from cell D2 to the southwest. Generated path frequency is overlaid onto the land bridge grid. Red shaded cells were alwayns visited by a portion of all simulated runs over all herd sizes and algorithms.

VI. Conclusions

Environmental tipping points can occur on a variety of scales. In this work artificial intelligence and virtual reality are used to study the impact that local behavioral changes are related to and reflect the impact of larger scale climatic change.

Several algorithmic models of herd movement were described and simulated using various parameter configurations. The result was to produce stress curves that showed how survivability was affected by herd size for each algorithm. From those curves a three-phase model of survivability was produced, and three representative herd sizes were selected for path optimization using an evolutionary algorithm, Cultural Algorithms.

Cultural Algorithms were used to produce optimal pathways for herds of different sizes. The results suggest that it was advantageous for herds of large size to expend their movement into the northern areas of the bridge. However, the reasons for the expansion were different for each model but the results were the same in principle. With this new knowledge in hand, the next step will be to investigate the impact of these changes on hunting strategies and hunter group movement in future work.

References

Amato, “Procedural Content Generation in the Game Industry,” in Game Dynamics, Cham: Springer International Publishing, 2017, pp. 15–25. doi: 10.1007/978-3-319-53088-8_2.

Culturally Responsive-Sustaining Education Framework, http://www.nysed.gov/crs/framework.

Drikx, S. Meijer, J. van der Velden, and A. Iosup, “Procedural content generation for games,” ACM Transactions on Multimedia Computing, Communications, and Applications, vol. 9, no. 1, pp. 1–22, Feb. 2013, doi: 10.1145/2422956.2422957.

Fogarty, J., Reynolds, R.G., and Salamayeh, A., and Palazzolo, T., “Serious Game Modeling of Caribou Behavior Across Lake Huron using Cultural Algorithms and Influence Maps”, Ed. John O’Shea and E. Sonnenberg, Archaeological, Ethnographic and Paleo-Environmental Perspectives on Caribou Hunting in the Great Lakes: Maps, Directions and Accommodations, University of Michigan Museum of Anthropology (2015).

Koenig, S., Fast Replanning for Navigation in Unknown Terrain, in IEEE Transactions on Robotics, vol. 21, no. 3, pp. 354-363, (2005)

Lewis, C.F.M., Understanding the Closed-Basin Phases (Lowstands) of the Laurentian Great Lakes and their Significance, *Geoscience Canada*, Vol. 43, pp. 179-197, (2016).

McLeod, S., & Shareski, D., Different schools for a different world. Solution Tree Press, Bloomington IN (2017)

Millington, I., AI for Games. Third edition. CRC Press, Boca Raton, FL (2019)

O’Shea, J, Lemke, A., Sonnenburg, E., Reynolds, R.G., and Abbot, B., “A 9,000 Year-Old Submerged Caribou Hunting Structure in Lake Huron”, Proceedings of the National Academy of Science (2014).

Reynolds, R. G., O’Shea, J., Che Z., Gawasmeh, Y., Meadows, G., and Fotouhi, F., “The Agile Design of Reality Game AI”, in Multi-Agent Applications with Evolutionary Computation and Biologically Inspired Technologies: Intelligent Techniques for Ubiquity and Optimization, IGI Global Press (2011).

Reynolds, R.G., Culture on the edge of chaos: Cultural algorithms and the foundations of social intelligence, 1st ed. Springer International Publishing Basel, Switzerland (2018).

Reynolds, R.G., Palazzolo, T., Lemke, A., O’Shea, J., Saad, S., Zhang, C., “Deepdive: The Use of Virtual Worlds to Create Ethnographies of Ancient Civilizations”, Human Computer Interaction International Conference, Washington D.C. (2021).

Reynolds, R.G., Cultural Algorithms: Tools for the Engineering of Social Intelligence into Complex Cultural Systems, Wiley-IEEE Press, Hoboken (2021)

The IB Learner Profile, https://www.ibo.org/globalassets/publications/recognition/learnerprofile-en.pdf.