Deep Social Learning in Dynamic Environments Using Subcultures and Auctions With Cultural Algorithms

Leonard Kinnaird-Heether
Department of Computer Science
Wayne State University
Detroit, MI, USA
lkinnaird@wayne.edu

Robert G. Reynolds, Senior Member IEEE
Department of Computer Science
Wayne State University
Detroit, MI, USA
reynolds@cs.wayne.edu

Abstract— **Cultural Algorithms have led to the development of many ways to distribute information within social networks. These mechanisms act by helping the system make decisions about how information is distributed through a population network, and thus are called distribution or decision mechanisms. Many distribution mechanisms have been developed using techniques from auction theory, game theory and various forms of voting construct. Here we discuss several methods of Knowledge distribution collectively called the auction distributions mechanisms and their performance is compared using dynamic complex real-valued functional landscapes. We perform this comparison with regards to robustness, how well the system finds solutions, and resilience, how well the system reacts to changes in the dynamics of the system. In this paper an additional Subcultured Distribution Mechanism is described that works to factor the knowledge distribution mechanism into subnetworks in order to support a "deep social learning" approach. The Subcultured Distribution Mechanism is compared with the results of each individual distribution mechanism without a subculture enhancement, when applied to a series of dynamic complex optimization problems of varying complexities. The results suggest that relatively simple mechanism such as Weighted Majority Wins and First Price Auction are sufficient for environments that exhibit low entropic levels of change such as in linear changing environments. For non-linearly changing environments, First Price Multi-round and English Auctions are most of effective on their own. The Subcultured Distribution Mechanism extension of these mechanisms was found to be best suited for complexities where the two distribution mechanisms had similar performances, and in the most chaotic environments where having multiple distribution mechanisms to choose from was advantageous.**

Keywords*- cultural algorithms, evolutionary algorithms, auction, subcultures, deep social learningsocially motivated, cultural engine, dynamic complex environments, sustainability, robustness, resilience and problem solving algorithms.*

I. INTRODUCTION

In the modern study of Artificial Intelligence, the development of algorithms designed to solve and work in complex systems has become increasingly important. A complex system can be defined as "…a dynamic network of many agents acting in parallel, constantly acting and reacting to what the other agents are doing..." [1]. One key aspect of the modern problem-solving system for dynamic environments is sustainability, or how well the system adapts to changes in the problem landscape over time. One can measure sustainability by two factors: robustness and resilience [2]. Robustness relates to how a system weathers the impact of a perturbation in the problem environment, much like how trees bend in the wind, but do not break. Resilience refers to how the system adapts to the environment in order to maintain or improve its performance over time after a major change event.

Cultural Algorithms are computational models of social evolution based upon principles of Cultural Evolution. A Cultural Algorithm is comprised of a Belief Space consisting of a network of active and passive knowledge sources and a Population Space of agents. The agents are connected via a social fabric over which information can be exchanged and used by agents. The knowledge sources in the Belief Space can compete or cooperate with each other in order to influence the decision making of agents in the Population Space. Likewise, the problem-solving experiences of agents in the Population Space are sent back to the Belief Space and used to update the knowledge sources there. It is a dual inheritance system in which both the Population and Belief spaces evolve in parallel. The Cultural Algorithm framework can be seen in Fig. 1.

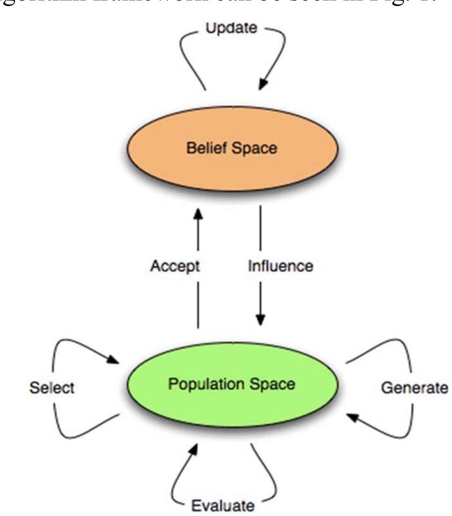

Fig. 1. The Cultural Algorithm framework

In the Cultural Algorithm [3], one of the major factors that affects the performance of the search process in a problem landscape is conflict resolution between competing Knowledge Sources in the Belief Space. This conflict resolution, the knowledge distribution mechanism, determines which of the Knowledge Sources is allowed to influence particular individuals in the networked Population Space. That is, if the individual is connected to neighbors who are influenced by conflicting knowledge sources, which Knowledge source should the individual select?

This conflict resolution process is managed using heuristics called **knowledge distribution mechanisms.** These mechanisms are modeled after real-world constructs such as voting [4], auctions [5] and game theoretic [6] approaches. This paper will discuss some of the auction-based distribution mechanism approaches, and how they compare to mechanisms developed in previous work. This paper will also discuss a new sub-cultured approach to distribution mechanisms that allows the Cultural Algorithm to learn to use a combination the different mechanisms for problems of a specific complexity, based on their previous performance in those situations. This is done by effectively developing subnetworks of information flow such that different mechanisms can operate on different aspects of a signal at a given subnetwork. A view of the spectrum of distribution mechanisms can be seen in Fig. 2. The idea of the figure is that as one moves counterclockwise from Random to Game Theoretic mechanisms, the fidelity of the signal provided needs to increase in order for the mechanism to be effective. For example, while there is no discernable signal for a random mechanism to work, a majority win or "wisdom of the crowds" mechanism works well when there is signal but some noise. Likewise, with increased fidelity the signal strength increased and some of what was previously noise becomes transformed into weights for different approaches. Auction mechanisms work well in situations where certain parts of the signal can be further quantified and made available to the competing knowledge sources in order to influence their respective bids.

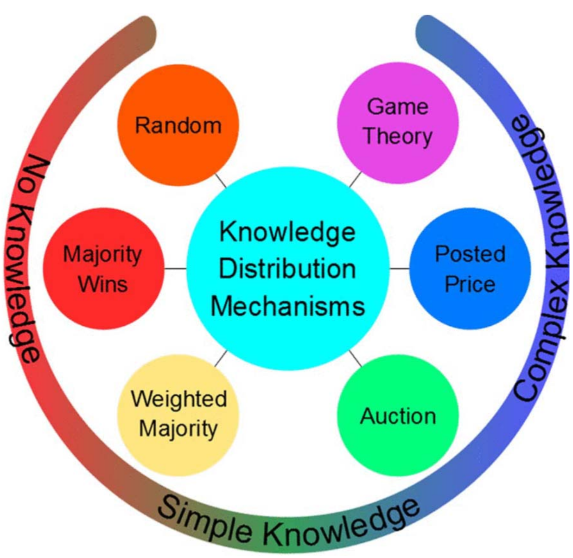

Fig. 2. The Spectrum of Knowledge Distribution Mechanisms

The **sub-cultured distribution mechanism** affords the system the opportunity for these mechanisms to operate on the signal in parallel just as subcultures do. It represents one mechanism for deep social learning. It is **social** in the sense that an individual can participate in more than one subnetwork and it supports "**deep social learning"** in the sense that the knowledge produced in one subnetwork can be transferred to another through the mutual participation of an individual. Thus, more than one distribution mechanism can be involved at different subnetworks. We view sub-cultures then as one avenue for the generation of "**deep social learning**" systems. Unlike "deep learning neural networks", deep social learning supports the involvement of an agent in more than one social network so that experiences in one network can be transferred to individuals that they are related to in other subnetworks. While other aspects of the subculture model have been addressed elsewhere [8], the focus here will be on a very fundamental example in order to demonstrate its ability to add value to the evolutionary learning process in even the simplest situation. Here, there is only one homogeneous network, but each knowledge source can be distributed by four different distribution mechanisms. Even in this simplest of cases it will be shown that a sub-cultured perspective can significantly improve a cultural systems' learning capacity over Cultural Algorithms without such a component.

This paper will also discuss how the Sub-cultured distribution Mechanism (SDM) system can be combined with the human-machine interface in the Cultural Engine [8]. By using the measures of robustness and resilience in the system as diagnostic signals, one can see how the various mechanisms perform under different complexities, thus informing a human-in-the-loop system where one can tailor the distribution mechanism use patterns to fit the problem complexity and allow the Cultural Algorithm Engine to fine tune its performance.

In the next section we will discuss the Cultural Algorithm and basic knowledge distribution mechanisms. The auction mechanisms are discussed in Section III and Section IV describes the new SDM system. Section V will discuss the Cones World environment, a framework for generating dynamic complex problem landscapes based upon the concept of entropy. Section VI will cover the experimental framework used for the results presented in this paper. Section VII will discuss our results on tests done using the individual auction mechanisms, and Section VIII will discuss results from tests done with the Subcultured Distribution Mechanism. Section IX will discuss how our results fit into the human-machine interface as part of the Cultural Engine and detail future work.

II. THE CULTURAL ALGORITHM

The Cultural Algorithm [9] is an approach to social evolution that employs a Belief Space or collection of Knowledge Sources (KS) that can be used to direct the problem-solving activities of individuals in the population. The Cultural Algorithm implements the five knowledge sources as part of the Belief Space. Each knowledge source can be viewed as a container of related knowledge sources. These five knowledge sources derive from the basic five categories of knowledge described in the field of Semiotics.

The basic Knowledge Source (KS) categories or containers are as follows. Normative knowledge recalls what is deemed to be an acceptable range of values for the solution to a given problem and uses this to guide the search for that solution. Situational knowledge is a record of important events, in the form of solutions that were particularly successful or unsuccessful. Domain knowledge uses specific information about the domain of the problem to influence the search. Historical knowledge keeps a temporal record of the search space. Topographical knowledge is spatial information about the topography of the search space *or* performance landscape. A pseudo-code representation of the Cultural Algorithm is shown in Fig. 3

```
Begin
    t = 0;
    initialize Bt, Pt
    repeat
        evaluate Pt {obj()}
        update(Bt, accept(Pt))
        generate(Pt, influence(Bt))
        t = t + 1
        select Pt from Pt – 1
    until (termination condition achieved)
End
```

Fig. 3. Cultural Algorithm Pseudo-code

The influence function determines how the knowledge sources align to influence the decisions made by individuals in the population. The influence function has two distinct steps: the direct influence step and the knowledge distribution or conflict resolution step. Together, they determine how the various heuristics are spread among the population of individuals in the population. The knowledge distribution mechanisms provide the means of resolving influence conflicts between knowledge sources over the right to influence an individual during the knowledge distribution step.

First, in the direct influence step every individual in the population is associated with an individual knowledge source as its **direct** influence [9]. This is done traditionally with a roulette wheel process, where the relative area of each KS on the wheel is a function of its relative performance against the other knowledge sources. This is done since here individuals do not have a prior memory of previous states. This is again consistent with a baseline scenario of minimal complexity for the agents in the network. Each individual agent is connected to a subset of others in the population via a network. That network can be homogeneous or heterogeneous in terms of the indegree and outdegree of its nodes. The underlying network is termed the Social Fabric since its structure can be changed over time based upon agent interactions.

Once the direct influence has been determined for each agent, the system checks to see if the direct influence for one agent conflicts with its adjacent neighbors in the network. If there is a conflict then a conflict resolution of knowledge distribution mechanism is called in to resolve the conflict. The knowledge distribution mechanisms used for this research are weighted majority (often called "the wisdom of the crowd"), and three auction mechanisms: first price, first price multi-round, and English auction. Of these methods, the least complex is the weighted majority, which assigns influence using a majority wins approach, as shown in Fig. 4.

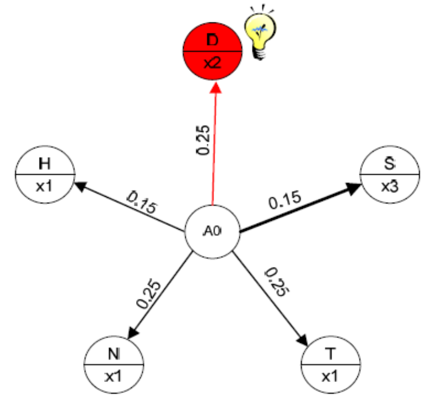

Fig. 4. The Weighted Majority distribution mechanism process.

In this figure, the top of each node represents a knowledge source category, while the bottom corresponds to the number of adjacent neighbors with that direct influence. The direct influence of A0 has been added into the count for its direct influence. The Domain knowledge source is assigned influence over an individual, even though it only influences two of that individuals' neighbors, compared to the Situational knowledge source that influences three. This is because the Domain knowledge source has a greater weight value of 0.25, than Situational which is only 0.15. In this work a homogeneous "four square" network topology is used for the population component since the sum of the neighbors plus the individual is an odd number, so it is more likely to produce a winner.

III. Auction Distribution Mechanisms

In this paper we will analyze auction mechanisms based on ascending-bid and first-price sealed bid auctions. These auctions were chosen because unlike the descending-bid auction, they use an ascending bid system that is more closely aligns with the fitness-maximizing nature of the Cultural Algorithm. The auction distribution mechanism [11] was originally designed to emulate the first-price sealed-bid auction. In this case, each knowledge source selected as a bidder spins their associated bidding wheel once to generate a bid, and the knowledge source with the highest generated bid will be declared the winner. In the case of a tie, the winner was chosen randomly between the knowledge sources that had submitted an equal value bid.

In addition to the first-price sealed-bid (FP) auction, we also use a variation on the original first-price sealed-bid solution, with the exception that instead of relying on a random selection to break ties, the solution uses a multi-round tiebreaking scheme, lending it the name First-Price Sealed-Bid Multiround (FPM). In the case of a tie, the knowledge sources who have submitted an equal value bid are allowed to spin their bidding wheels again and thus submit a new bid. This repeats in the case of further ties, until one knowledge source submits a singularly higher bid than the other remaining knowledge sources and is thus declared the winner.

The other new auction distribution mechanism that has been added is designed to emulate the classic English auction. In order to represent the unordered nature of the bidding process inherent to the English auction, after the participating knowledge sources are defined, an initial bidder is chosen at random. The bidder then spins its bidding wheel to generate a bid. Every other competing knowledge source is then allowed to bid in response until a higher bid is generated. This process repeats until no higher bids are generated. Then the knowledge source with the current highest bid is declared the winner.

IV. Subcultured Distribution Mechanisms

A complex system can be described at several levels of functional granularity. One vehicle to support parallel learning and adaption is the concept of subcultures. Subcultures can be viewed as networks that connect a subset of individuals that extract and process a subset of information and pass that information to other subnetworks through shared individuals [7]. One of the principle functions of a social system is to provide a sustainable environment for its participants. Sustainability results from the ability of the system to withstand and adapt to changes in a dynamic environment. These qualities have been discussed earlier in terms of the factors of resilience and robustness. The changes in a dynamic environment can be produced by a variety of factors. Subcultures allow the system to operate on an incoming information in different ways. In this case a sub-culture is a specialized network that serves to transmit a specialized amount of knowledge or signal to a subset of the population. These subcultures are connected together by shared sets of individuals. In this research, subcultures are a device that enables deep social learning in a population through interacting subnetworks in the Cultural Algorithm.

However, the knowledge about the social fabric's topologies in the Population Space is not currently visible to the Belief Space. Thus, by collecting the individuals' experiences in topologies and knowledge sources, we expect to see subcultures start to evolve in the Population Space. The Subcultured Distribution Mechanism as currently implemented follows a roulette wheel-based approach to choose between decision mechanisms, and it reconfigures the selection process in order to prevent knowledge source dominance. It is therefore able to identify certain knowledge source distribution mechanism that are relatively effective during a certain portion of the systems evolution.

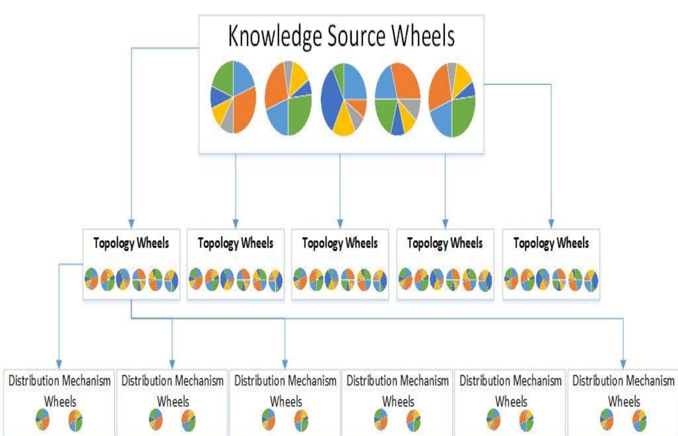

Fig. 5. The Subcultured Distribution Mechanism

In this approach, the subculture controller maintains a selection wheel for each knowledge source, shown in phase one in Fig. 5. That wheel contains an entry for each of the topologies being used, shown in phase two. The subculture controller maintains a selection wheel for each topology/knowledge source pair. This wheel contains an entry for each of the distribution mechanisms being used, shown in phase three. This second layer of selection wheels is the key component of the Sub-cultured Distribution Mechanism. Each selection entry is based on the performance of that selected knowledge source in the previous generations. This allows the system to maximize particular entries that might be better suited for the particular problem at hand. An example of this approach is given in Figure 5.

Once the topology is selected in phase two and the distribution mechanism is selected in phase three, the Cultural Algorithm proceeds through the standard process of a generation, using the topologies and distribution mechanisms selected by the Sub-cultured Distribution Mechanism whenever required. After the generation has been completed, the selection wheels for the next generation are calculated. In the original subcultures model, each selection wheel entry received a portion of the wheel that is equated to the proportion of individual fitness values associated with the selection entry to the overall fitness. However, this method lets entries that do well early on to dominate the selection wheels, producing a positive feedback loop that worked to marginalize the other knowledge sources. This issue was corrected in the Sub-cultured Distribution Mechanism by simplifying the method by which the selection wheels were being altered.

The Subcultured Distribution Mechanism maintains a running average of the performance in terms of fitness generated by individual population or knowledge source in the purview of each network topology and distribution mechanism combination. If at any given generation i, the performance is better than the performance of that topology/distribution mechanism combination through generation $i-1$, the likelihood of use is incremented by a value v, and vice-versa. If the likelihood of use for a topology or distribution mechanism has reached a predefined upper or lower bound, then no change is made. By keeping the increment value v small and introducing the upper and lower bounds, we are able to keep any one method from dominating the system, or dropping out completely.

V. Cones World

Cultural Algorithms are a data-driven evolutionary hyper-heuristic algorithm that, like human cultures, allow the use of various knowledge sources to deal with dynamic changes in problem environments over time. The current Cultural Algorithm toolkit (CAT) contains both static and dynamic problem generators. The static component consists of engineering benchmark problems and the dynamic component is a dynamic landscape generator based on the work of Langton [12].

What evolved into the Cones World model was originally developed by De Jong and Morrison [13] as a dynamic landscape generator. It was modified by Reynolds and Ali [14] in order to examine the ability of evolutionary problem-solving approaches to solve randomly generated problems of arbitrary complexity. This was proposed as an alternative to the traditional approach of comparing algorithms on a small set of benchmark problems. It was felt that by focusing on a small set of problems, investigators often sacrificed algorithm generality in order to produce good results on the specific problems.

The Cones World Generator produces a problem landscape, in which a field of resource cones of different heights and different slopes that are randomly scattered in a multi-dimensional landscape. To this end, the Cones World Generator creates a dynamic environment in two steps: first by specifying a baseline static landscape of the desired morphological complexity, then by incrementally adding the desired dynamics. The equation that all Cones World environments are derived from is given by the equation in Fig. 6:

Eqn 1

$$f(\langle x_1, x_2, \dots, x_n \rangle) = \max_{j=1,\mathrm{k}} \left(H_j - R_j \cdot \sqrt{\sum_{i=1}^{n} \left(x_i - C_{\mathrm{j,i}}\right)^2} \right)$$

Where:
k : the number of cones,
n : the dimensionality,
H_j: height of cone j,
R_j : slope of cone j, and
$C_{j,i}$: coordinate of cone j in dimension i.
The values for each cone (H_j, R_j, and $C_{j,i}$) are randomly assigned based on the following user specified ranges:
$H_j \in (H_{base}, H_{base} + H_{range})$
$R_j \in (R_{base}, R_{base} + R_{range})$
$C_{j,i} \in (-1, 1)$

Fig. 6. Equation for Generating a Cones World Environment.

The max function in equation 1 given in Fig. 6. is used to blend each of the independently specified cones together so that if, for example, two cones overlap, the height at a point is the height of the cone with the largest value at that point. An example landscape is shown in Fig. 7.

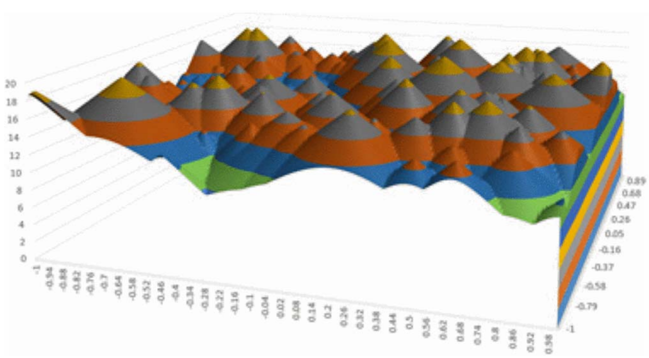

Fig. 7. An Example of a Cones World Landscape

This cone landscape generator function can be specified for any number of dimensions, and once the base static landscape is generated, the second step is to specify the environment dynamics. For each cone j and dimension i, a parameter can be changed individually and independently that will control the complexity of the landscape. The value $C_{j,I}$ controls the position of the cone in i-dimensional space, H_j the height, and R_j the slope. By applying changes to these parameters, a single time, a static complex landscape can be generated. For instance, one could choose to only apply changes to the positions of the cones, or to cone height, or cone slope. In addition, the Cones World generator can use this functionality to alter the problem landscape while the Cultural Algorithm system is analyzing it, thus creating a dynamic environment that can be changed to allow us to test the resilience and robustness of the current Cultural Algorithm configuration. In this case robustness refers to the ability of the system to regain or improve its performance in the wake of a series of small change events such as blows in a boxing match, while resilience corresponds to the system's ability to bounce back after a major change event such as a knockdown in a boxing match..

By changing the environment dynamics in a specific way, we can generate problem landscapes that represent complex systems with varying levels of dynamic components with which to test various Cultural Algorithm configurations. This enables us to evaluate our model in a more flexible and systematic way. It also is a reasonable facsimile of how resources are spread out within natural environments. From an information theoretic point of view, the problem environment carries a signal of a certain complexity of information that can be represented in terms of its entropy. This entropy value relates to how much information cells in a cell space need to know about their neighbor cells. In our cones world this means that as the distribution of cones becomes more complex, agents will need more information to find their way to the optimum.

The complexity of the environment is based on the work of Langton [12] that describes the A value as it relates to entropy of the system. The A values can range between 0 (for static) and 4.0 (for dynamic), and as the value increases the system becomes more unpredictable. Langton demonstrated this in terms of rules needed by a Cellular Automata to solve problems at different points. A graph of logistic functions described by Langton's work is provided in Fig. 8. This shows the values of Y that are generated for each iteration of the logistic function, given values of A between 1.0 and 4.0. A = 0.0 represents the static case and is not shown here. For values of A=3.0 or more the system switches from linear to nonlinear in its behavior. Previous work ha shown that the wisdom of the crowds does well in linearly changing environments but begins to degrade in non-linear ones. This is where one expects to see mechanisms such as auction perform better since there is more signal to extract and make use of.

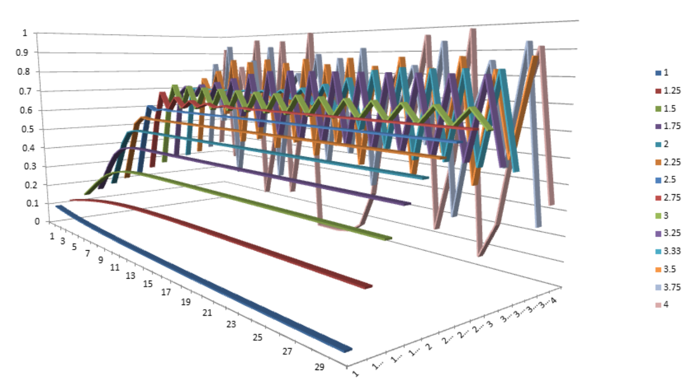

Fig. 8. Logistic Function for Selected Complexity A-values

In particular, Langton confirmed that the amount of mutual information that cells in the space needed to know about their neighbors increased as the entropy or amount of information in the landscape increased. Based upon that he established several basic computational classes as shown in Table I.

TABLE I. TABLE OF A-VALUE CLASSIFICATION

Classification	Range	Description
Static	A = 0.0	Not covered in this research.
Linear	0.0 < A < 3.0	For problems of low entropy, a fixed set of rules can be given to each cell in order to allow them to exchange the information needed to solve the problem.
Non-Linear	3.0 ≤ A ≤ 3.8	For problems of this nature the cells need to switch from one set of rules to another depending on the number of bifurcations.
Chaotic	*A ≥ 3.9*	Problems for which the number of bifurcations is so large the system is inherently chaotic.

VI. EXPERIMENTAL FRAMEWORK

The experiments here were developed to collect results of the Cultural Algorithm toolkit (CAT) system on multiple dynamic Cones World landscapes. These tests were used to compare the weighted majority win or "wisdom of the crowds" (WMW) distribution mechanism with the auction distribution mechanisms English, First Price and First Price Multi-round on their own, without the subculture distribution model. In addition, statistical tests were also performed in order to compare the aforementioned results. Similar statistical tests were made against the enhanced subcultures system model as well.

The experiments were developed to collect results of the CAT system on multiple dynamic Cones World landscapes. These tests were used to compare the weighted majority win (WMW) distribution mechanism with the auction distribution mechanisms English, First Price and First Price Multiround on their own, against the Sub-cultured Distribution Mechanism (SDM). In addition, statistical tests were also performed in order to compare these results with results from testing the SDM. In these tests the goal is to investigate how distribution mechanisms and the Subcultured Distribution Mechanism perform under various complexities, from the linear to the non-linear and chaotic. We also want to determine how the system behaves during and after the landscape shifts. In particular, we seek to analyze how long it takes for the system to find a new optimum solution after a major change in performance landscapes ,and whether that length of time to do so is reduced as the test progresses, and how the system reacts to sudden but small changes within a given landscape. The former is representative of resilience, as was described previously, and the latter a measure of robustness. These two metrics are critical to our discussion of sustainability as it pertains to Cultural Algorithms and dynamic complex environments. Specifically, we are going to test the system against the following two points:

In order to demonstrate how the modifications to the Cultural Algorithm detailed in this document address these issues, a series of tests were conducted. Each test consisted of running the CAT system on a given Cones World landscape for 1000 generations. Then, a change in the Cones World based upon the current A values is made. A sequence of 50 landscape window changes was produced with the system having 1000 generations to explore each. The same sequence was used for each test to afford a fair comparison between approaches. At the end of each window the Cones World configuration is adjusted but the current state of the population and belief space is not. Therefore, the CA will need to adjust its performance in order to track this change. System resilience will improve if it is able to track these changes over time.

If the system finds the maximum point in the landscape (within a threshold value, in this case 0.0001) before the 1000-generation limit, we say that the solution is found. If the solution is not found, the dynamic change event is called. Once the generation limit is reached, the Cones World generator is called and the characteristic A-values are applied to the landscape. In these experiments only changes in the location of the cones in the landscape are made. This change process is repeated 50 times for each mechanism, giving us a total of 50000 generations over a dynamic landscape for each run of a mechanism. Every mechanism was run 50 times for the window sequence. These experimental parameters are similar to those used in previous work [15], however the threshold for solution value was lowered from 0.001 to 0.0001 to provide a slightly increased complexity. We fix the network topology using only the Square topology. This topology has been shown in the past [5] to provide a balance between the limited connectivity of the L-Best network and the full connectivity of the Global network. Fixing the topology will allow the contrast of the individual mechanism results with the Sub-cultured Distribution Mechanism results without having to consider whether topology differences are having an effect on the results and fix the focus on the differences in the distribution mechanisms.

One of the major goals of this research is to determine how the different distribution mechanisms perform at varying levels of complexity. To this end a wide range of complexity values have been chosen to test the different complexity classes. These values are A = 1.4, 1.8, 2.2, 2.6, 3.1, 3.2, 3.3, 3.4, 3.6, 3.7, 3.8, and 3.9. Each test used the Square network only in order to provide clarity when comparing each of the distribution mechanisms.

VII. RESULTS COMPARING THE INDIVIDUAL DISTRIBUTION MECHANISMS

The distribution mechanisms examined here exist over the spectrum of knowledge distribution mechanisms described earlier. Weighted Majority Wins requires simple knowledge to function and is expected to perform better on less complex landscapes where the entropy is low and the amount of information required to find the solution is less. The auction mechanisms as a whole are more complex than Weighted Majority Wins, and thus are better able to function when the knowledge requirements available to find the solution are more quantitative [15]. Additionally, each type of auction mechanism individually has its own knowledge requirement. The First Price auction is the least complex, consisting of only a single bidding round, and thus is limited in its knowledge distribution capacity with respect to the other auction mechanisms. However, it is still more complex than other mechanisms such as Weighted Majority Wins. The First Price Multi-round auction increases the complexity by adding additional tie-breaking rounds, and the English auction rounds out the auction mechanisms as the most complex.

Table II shows the average time to completion over all 50 landscapes tests in all 50 runs for each of the distribution mechanisms tested, Weighted Majority Wins, English Auction, First Price Auction and First Price Multiround Auction. These results are organized by complexity and the mechanism with statistically most significant performer highlighted in green for each category. The overall average is shown in the last row. The distribution mechanisms are ordered from left to right in terms of the amount of knowledge or information within the system required for each distribution mechanism to function successfully, with the least amount being on the left.

TABLE II. AVERAGE NUMBER OF GENERATIONS TO COMPLETION FOR WEIGHTED MAJORITY WINS AND ENGLISH, FIRST PRICE AND FIRST PRICE MULTIROUND AUCTION, BY COMPLEXITY

Distribution Mechanism/ Complexity	Weighted Majority Wins	First Price	First Price Multiround	English
1.4	80.11181	83.42528	86.84806	91.56074
1.8	98.94212	99.85733	106.4458	121.2768
2.2	117.2493	116.2966	118.1474	121.1711
2.6	129.9544	113.1501	123.9874	126.7178
3.1	142.8914	133.6888	137.5147	135.1407
3.2	139.8533	139.9429	130.0718	134.0834
3.3	151.5514	142.9665	145.8261	126.0652
3.4	143.79	134.7416	133.2854	128.154
3.6	152.077	148.8036	144.1027	143.0955
3.7	160.6921	151.1682	137.9138	141.1885
3.8	153.12	148.1671	133.7517	130.5032
3.9	142.7504	148.8467	155.9442	126.3248
Overall	134.4153	130.0879	129.4866	127.1068

The complexities in the range of 1.4 to 2.6 lie within the Linear computational class. These complexities have a relatively low entropy and thus all distribution mechanisms should work well, as is the case. As expected, the First Price auction and Weighted Majority Wins distribution mechanisms perform best in this range. These are the least complex of the distribution mechanisms tested and are well suited to perform well in the Linear class of landscapes because the amount of information required for them to function well is less than that of the others. The more complex mechanisms, First Price Multiround and English Auction, perform well, but their extra computational complexity works against them, leading to slower information discovery. Since the information moving within the networks in the Population Space and into the Belief Space is relatively predictable, the resilience provided by the complex distribution mechanisms is not needed.

The complexities in the range of 3.0 to 3.8 are of the Non-Linear complexity class. These complexities produce more complex environments than the Linear class, but less complex than the Chaotic class. Although deterministic, they are the product of the convolution of small adjustment. The increased complexity requires more information about the problem landscape to travel within the system, meaning that the more complex distribution mechanisms should outperform the less complex ones. Here we can see that the First Price Auction performs best only in the lowest complexity in this range 3.1. As the complexity increases, the more complex distribution mechanisms, First Price Multi-round and English Auction, alternate for the title of best performer on average. First Price remains in contention until the complexity reaches 3.4. It then falls off. Weighted Majority Wins is the worst performer for these nonlinear environments in almost every case.

The complexity A = 3.9 is the last complexity value we tested and lies within the Chaotic complexity class. Complexities in this class produce the most complex environments and require the most information about the problem landscape to within the system. In this case, the English auction, the most complex distribution mechanism tested here, vastly outperforms all of the other mechanisms.

When examining the cumulative average time to solution over all complexities tested, the English auction provides the best result. As we have observed, each of the distribution mechanisms had regions within the complexity spectrum where they were successful. This finding will be useful when discussing how the Sub-cultured Distribution Mechanism is able to switch multiple individual distribution mechanisms. This should enable the Cultural Algorithm to pick and choose when to use a given knowledge source distribution mechanism combination, allowing it to maximize their potential based on the complexity of the problem.

VIII. RESULTS COMPARING THE SUBCULTURED DISTRIBUTION MECHANISM

The Sub-cultured Distribution Mechanism employs all of the distribution mechanisms, Weighted Majority Wins and First Price, First Price Multi-round and English auction tested in the previous section. Each distribution mechanism is initially given an equal likelihood of being selected ,and this value is increased or decreased based on the mechanisms' performance. The average time to solution analysis of the individual distribution mechanisms given in Section VII showed that there is a clear relationship between the complexity of the distribution mechanism and complexity of problem landscape. In addition, there are several A-value ranges where one single distribution mechanism does not necessarily have a distinguishably lower average time to solution than all of the others. These areas occur as the complexity increases and the best performing distribution mechanism transitions from one to another mechanism. This scenario suggest that there is an opportunity for several mechanisms to work together in these transitional stages.

Table III shows the average number of generations needed to complete the solution over all complexities for each individual mechanism including the Subcultured Distribution Mechanism. As in Table II, the green highlights indicate the lowest average time to solution of all the individual distribution mechanisms for a given complexity class. The black-border highlights indicate those A-values where the lowest average time to solution was found when testing the distribution mechanisms alone. Each of these represents a complexity class where that performance was improved upon by the Sub-cultured approach. Notice that the majority of the classes, are clearly transitional in nature. Classes 3.1 and 3.2 represent the result of a gradual phase shift from linear to non-linear. Likewise, classes 3.6 and 3.7 correspond to a transition from non-linear to chaotic.

TABLE III. AVERAGE NUMBER OF GENERATIONS TO COMPLETION FOR WEIGHT MAJORITY WINS. ENGLISH, FIRST PRICE AND FIRST PRICE MULTIROUND AUCTION AND SUBCULTURED DISTRIBUTION MECHANISM, BY COMPLEXITY

DM/ Complexity	**WMW**	**First Price**	**FPM**	**English**	**Subcultured Distribution Mechanism**
1.4	80.11181	83.42528	86.84806	91.56074	80.93344
1.8	98.94212	99.85733	106.4458	121.2768	104.8647
2.2	117.2493	116.2966	118.1474	121.1711	114.9718
2.6	129.9544	113.1501	123.9874	126.7178	129.067
3.1	142.8914	133.6888	137.5147	135.1407	132.7802
3.2	139.8533	139.9429	130.0718	134.0834	127.3045
3.3	151.5514	142.9665	145.8261	126.0652	140.723
3.4	143.79	134.7416	133.2854	128.154	131.62
3.6	152.077	148.8036	144.1027	143.0955	132.5028
3.7	160.6921	151.1682	137.9138	141.1885	135.7517
3.8	153.12	148.1671	133.7517	130.5032	137.0336
3.9	142.7504	148.8467	155.9442	126.3248	134.0947
Overall	134.4153	130.0879	129.4866	127.1068	125.1373

In the Linear Complexity Class, both Weighted Majority and the First Price auction maintained their dominance in three of the four instances tested, but the Subcultured Distribution Mechanism slightly outperforms them both for A=2.6. As the environment's complexity increases in their A-value, neither First Price or First Price Multi-round are dominant in any of the A-value classes as they were previously. Instead, the Subcultured Distribution Mechanism had the lowest average time to solution in A = 3.1, A=3.2, A=3.6, and A=3.7. In these A-value regions, the Subcultured Distribution Mechanism takes advantage of the multiple distribution mechanisms under its control and learns to combine them to improve its performance over each individual. Now that the Subcultured Distribution Mechanism has information about the relationship between knowledge sources and distribution mechanisms it requires fewer generations to solve the problem in those instances.

We also note that English Auction and Sub-cultured Distribution Mechanism are the dominant mechanisms in all of the non-linear instances tested from A=3.1 through A= 3.9, with each taking four complexity classes apiece. There is a clear complementarity between those two approaches, with each of those two mechanisms dominated for two complexity instances, followed by the other in the subsequent two. Note that the English Auction was the best performer of the four auction categories in terms of its overall average performance across all complexity classes. However, the Sub-cultured approach is the clear winner since it is able to transition from one type to another as the complexity changes.

IX. CONCLUSION AND FUTURE WORK

In this paper we investigated Cultural Algorithms as a means for problem solving in dynamic complex environments and the concept of sustainability in complex systems, or how well the system adapts to changes in its problem landscape over time. We first tested each of four different auction mechanisms on their own against a spectrum of complexity classes ranging from linear to chaotic. It was clear that each had its own "sweet spot" in terms of the complexity classes tested. The "wisdom of the crowds" approach did will in environment with small linear changes over time (A=1.4 and 1.8). However, once the slope of linear change became larger, there was sufficient additional information to support the simplest auction mechanism, First Price. It dominated during the transition from linear to non-linear complexity (A=2.2, 2.6, and 3.1). However, once solidly into the non-linear territory it was less able to process sufficient information to track the environment in that phase. First Price Multi-round then took over for A=3.2. After that, the English dominated through to the chaotic phase although First Price Multi-round rebounded during the transition from Non-linear to Chaotic for A=3.7.

It was therefore clear that each mechanism had a complexity range over which it excelled based upon the difference in the knowledge needed to successfully track the movement of the optimal cone in the changing landscape. Each of these auction mechanisms had an informational "niche" in which it was most effective. These differences were the basis for the success of the Sub-cultured mechanisms. As the complexity increased it was able to combine predecessor and successor auctions together to form an informational "biome". The nature of these biomes then changes with changes in the informational content of the signal. Taken together these combinations were better able to track the systems changes and predict new ones than either component could do alone. The first biome was the "Weighted Majority and First Price" combination. The second was the "First Price and First Price multi-round" which was succeeded by the "First Price multi-round-English auction combination. While the other auction mechanisms contributed these pairings were the most dominant in each biome.

This work has demonstrated the power of Subcultures to support deep social learning even when there are few options to consider. Here, we only investigated the process of distribution mechanism selection as a baseline approach. This provides us with an understanding of why subcultures become more powerful tools for social learning when additional options for subcultural configuration such as subnetwork topologies and differences in information content come into play as has been suggested in previous work [9]. If one sees Subcultures as a mechanism for deep social learning in the Cultural Engine, the key will be what configurational choices are made available for subcultures to exploit. We suggest that the supervision of the system by a "human in the loop" who is able to inject or remove from the social system, when appropriate, based upon the values for Cultural Engine performance metrics as discussed in [9]. That will be the focus of future work.

REFERENCES

[1] M. M. Waldrop, *Complexity: The emerging science at the edge of order and chaos.* Simon and Schuster. 1993.

[2] J. Husdal, "Reliability and vulnerability versus cost and benefits." *Proc. 2nd Int. Symp. Transportation Network Reliability (INSTR)*. 2004, pp. 180-186.

[3] R. G. Reynolds, "On modeling the evolution of hunter-gatherer decision-making systems" *Geographical Analysis* vol. 10 no. 1. 1978, pp. 31-46.

[4] X. Che, M. Ali and R. G. Reynolds, "Robust evolution optimization at the edge of chaos: Commercialization of culture algorithms," *IEEE Congress on Evolutionary Computation*, Barcelona, 2010, pp. 1-8.

[5] R. G. Reynolds and L. Kinnaird-Heether, "Optimization problem solving with auctions in Cultural Algorithms," *Memetic Computing*, vol. 5, no. 2. 2013, pp. 83–94.

[6] F. Waris and R. G. Reynolds, "Optimizing AI Pipelines: A Game-Theoretic Cultural Algorithms Approach," *2018 IEEE Congress on Evolutionary Computation (CEC)*, Rio de Janeiro, 2018, pp. 1-10.

[7] R. G. Reynolds, Y. Gawasmeh & A. Salaymeh, "The Impact of Subcultures in Cultural Algorithm Problem Solving." *Computational Intelligence, 2015 IEEE Symposium Series on*. 2015.

[8] R. G. Reynolds and L. Kinnaird-Heether, "Population mechanics and cultural algorithms in the development of a cultural engine," *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, Honolulu, HI, 2017, pp. 1-8.

[9] R. G. Reynolds, "Cultural Algorithm Framework." *Culture on the Edge of Chaos*. Springer, Cham, 2018. pp. 21.

[10] R. G. Reynolds, M. Z. Ali, "Computing with the Social Fabric: The Evolution of Social Intelligence within a Cultural Framework", *IEEE Computational Intelligence Magazine*, vol. 3, no. 1. 2008, pp. 18-30.

[11] R. G. Reynolds and L. Kinnaird-Heether, "Problem solving using social networks in Cultural Algorithms with auctions," *2017 IEEE Congress on Evolutionary Computation (CEC)*, San Sebastian, 2017, pp. 2714-2721.

[12] R. G. Reynolds, M. Z. Ali, "Exploring knowledge and population swarms via an agent-based Cultural Algorithms Simulation Toolkit (CAT)." *Evolutionary Computation, 2007 IEEE Congress on*. 2007, pp. 2711-2718.

[13] C. Langton, "Life at the edge of chaos", *Artificial life II*, vol. 10. 1992, pp. 41-91.

[14] R. Morrison, K. De Jong, "A test problem generator for nonstationary environments", *Proc. Congr. Evol. Comput.* 1999, pp. 2047-2053.

[15] R. G. Reynolds & L. Kinnaird-Heether, "Diversity as a Necessity for Sustainability in Cultural Systems: Collective Problem-Solving in Cultural Algorithms." *Humanized Computing and Communication (HCC) 2019 IEEE International Conference on*. 2019, pp. 24-30.