

THE EL FAROL BAR PROBLEM AND COMPUTATIONAL EFFORT: WHY PEOPLE FAIL TO USE BARS EFFICIENTLY

WILLIAM RAND*, Northwestern University, Evanston, IL
FORREST STONEDAHL, Northwestern University, Evanston, IL

ABSTRACT

Does how much an agent thinks about its own actions affect the global properties of a system? We use the El Farol Bar Problem to investigate this question. In this model, the El Farol Bar represents a scarce resource. Does the amount of computational ability that agents possess affect resource utilization? For instance, if agents attend the bar randomly on average 50 people will go to the bar. On the other hand, if agents act as neoclassical economics suggest, its not clear what the average attendance at the bar will be, but in this paper we will argue that it will also be near 50. In Arthur's original model, he showed, using a simulation involving an ecology of strategies, that the average attendance of the bar converged to 60. Fogel et al. gave their agents more computational power and let them use a evolutionary algorithm; they showed that the average attendance at the bar was 56, not 60. If we examine these four results of (1) random agents, (2) perfect agents, (3) Arthur's agents, and (4) Fogel et al.'s agents, we can ask whether there is a relationship between computational effort and attendance at the bar (e.g. the utilization of a public resource). To investigate this question we look at a model where we can control the computational power that each agent has to predict the attendance each week.

Keywords: machine learning, agent-based modeling, El Farol Bar Problem, genetic algorithms

INTRODUCTION

Truly adaptive agents are one of the promises of agent-based modeling, but they are rarely used. This is particularly surprising since adaptation is one of the advantages that many people list as a reason to use ABM instead of other modeling techniques. When Holland (1995) discussed complex adaptive systems (CAS) and their relationship to ABM in *Hidden Order*, he devoted an entire chapter to adaptive agents, and specifically mentioned internal models as one of the mechanisms that define a CAS, and one of the most classic agent-based models, the El Farol Bar Problem, utilized adaptive agents.

In 1994, Arthur posed a problem he called the El Farol Bar Problem. The El Farol Bar is in Santa Fe, New Mexico and on Thursday nights it plays Irish music. There are 100 people in Santa Fe who like Irish music, but if more than 60 of them attend the bar then the bar is too crowded and no one enjoys the bar. If everyone attends the bar randomly, i.e. each Thursday they flip a coin to decide if they should attend, then the bar will be underutilized since on average 50 people will go to the bar. In other words, if the agents spend no computational effort, the bar is underutilized. On the other hand, we could consider what would happen if the agents act as neoclassical economics suggest and each agent does their best to predict the attendance of the bar (i.e. uses an infinite amount of computational effort). In this case, there seem to be two possible results, either (1) each agent predicts exactly the same attendance at the bar, in which case every agent will either go to the bar or stay home, or (2) since there is an infinite number of

* *Corresponding author address:* William Rand, Northwestern Institute on Complex Systems, 600 Foster Street, Evanston, IL 60208-4057; e-mail: wrand@northwestern.edu.

ways to predict the next number in a finite sequence each agent will have a different prediction; assuming that half of these predictions are that the bar will be crowded and half that it will not be crowded, then half of the ‘neoclassical’ agents will go to the bar, and half will stay home. In either case the average attendance at the bar will be 50, and the bar will be underutilized.

Arthur (1994) suggested a third model. He gave each agent a ‘bag of strategies’ that they could use to predict the bar attendance and every week the agents used the strategy that would have worked the best had they used it in the previous weeks. Arthur showed, using a simulation, that the average attendance of the bar using this ‘ecology’ of strategies wound up being 60 and thus under this model the bar was maximally used. At the end of his paper, Arthur speculates that if he had used an evolutionary algorithm instead of the ‘bag of strategies’ technique that the results would have been similar.

Fogel et al. (1999) decided to take up Arthur’s challenge and re-ran the El Farol simulation but allowed each agent to use an evolutionary algorithm with 10 strategies that evolved over 10 generations each week. Fogel et al. showed that in their model the average attendance at the bar was 56, not 60. The agents were doing better than random, but the attendance was not at 60, as Arthur had originally suggested.

If we examine these four results of (1) random agents, (2) perfect agents, (3) Arthur’s agents, and (4) Fogel et al.’s agents, we see that there is a different amount of computational power being employed and a different average attendance at the bar. Is there a relationship between computational effort and attendance at the bar (e.g. the utilization of a public resource)? We will investigate this question in this paper. We begin by examining some related background to the question at hand. From there we pose a possible hypothesis and postulate what it would entail. We then build a model to test this hypothesis, where we can control the number of evaluations that each agent carries out each week (i.e., the amount of computational power given to each agent). We explain how this model fits within our general framework for agent-based modeling and machine learning, and we discuss the design of this model. We then present the results of an experiment where we varied the computational effort of each agent. Finally we conclude by discussing these results within the larger context of adaptive agents, and the trade-off between computational effort and resource utilization.

BACKGROUND

Arthur’s original paper (1994) was more concerned with critiquing neoclassical economics than it was with investigating the particular properties of the El Farol Bar Problem. However, his research still represents one data point in our investigation into how computational power affects resource utilization. Fogel et al.’s extension (1999) of Arthur’s work was more concerned with questioning the stability of the attendance at the bar and how randomness and evolution of strategies affected the long-term dynamics of the El Farol Bar Problem. Still this work has also started to answer the question of how computational ability affects resource utilization. Wolpert et al. (2000) examined how to automatically configure agents to best utilize a bar. Whereas we are interested in how computational capabilities affect the overall performance of the system when the agents do not care about the overall resource utilization, Wolpert et al. take a more engineering approach to the question at hand, and design a system that attempts to optimally utilize the bar. There has been other work on the El Farol Bar Problem (Edmonds, 1999) and its refined version, The Minority Game (Challet et al., 2005), but the relationship between computational power and resource utilization is rarely investigated.

HYPOTHESIS

Our relationship of concern is the correlation between computational power and resource utilization. Though the connection between these two variables may be impossible to discover in general, we can begin by investigating it within the scope of the El Farol Bar Problem. As we laid out in the introduction, there are four data points that we already have to help us investigate this relationship. The first is random agents, which result in an average attendance of 50. The second is neoclassical agents, which we argue will result in an average attendance of 50. The third and fourth data points are Arthur (1994) with an attendance of 60 on average, and Fogel et al. (1999) with an attendance of 56 on average. How exactly to relate these to computational power is difficult, but clearly random agents possess the least computational power, since they do no computation except flipping a coin, and neoclassical agents possess the most computational power, since they are assumed to have infinite resources. The results from Arthur and Fogel et al. both fall somewhere between these two extremes. If we think of computational resources as the number of evaluations that we allow each agent to perform on its pool of strategies with the current history of attendance then we can actually quantify this resource. In Fogel et al.'s case this is easy. Each agent runs an evolutionary algorithm each week in which it evaluates 20 strategies (10 parent strategies and 10 child strategies) for 10 generations, resulting in 200 evaluations. In Arthur's case this is more difficult since he does not specify how many strategies each agent has in his original paper, but he lists as example numbers of strategies, 6, 12, and 23. Since Arthur's model does not create new strategies and just evaluates the extant strategies, then we can use 6, 12, or 23 as an approximation to the number of evaluations Arthur's agents carry out each week. Regardless Arthur's agents use less computational resources each week than Fogel et al.'s but more resources than the random strategy. We can graph these (rough) data points on a figurative plot (see Figure 1).

In Figure 1, there is a line representing a possible relationship between the number of strategy evaluations and the average attendance at the bar. This line is a hypothesis, but it is a reasonable hypothesis. The general reasoning is that if the agents have too little computational power then their behavior is essentially random. However, if the agents have too much computational power then their predictions start to resemble each other. Since the number of data points that the agent is using to evaluate each strategy is small (usually around 10), there is only so much data that the agents possess. As a result, after a certain point additional refinement of strategies does not result in an improved prediction. If agents were able to remember which strategy they had used in the past and make sure and choose a different strategy then the results might be different, but in the current model if the agents have found a strategy which correctly predicts the previous attendance then they will stay with it regardless if the same strategy failed them in the past. As a result, in the end the strategies of all agents will start to look similar. The more similar the agents' strategies look, the more likely all of the agents are to take the same action. In the extreme, if all agents take the same action then each week they will all go to the bar or all stay home, assuming they stay home or go to the bar with a uniform probability, this will result in an average attendance at the bar of 50.

On the other hand, if the agents are boundedly rational, and only possess a limited amount of computational power their strategies will likely be very different from each other. This will result in each agent choosing a strategy that works fairly well, but is also likely to be different than the other agents. This will create the 'ecology' of strategies that Arthur discusses. Given limited computational power it is unlikely that the agents will find an optimal strategy for the past n weeks and thus they will all converge to suboptimal solutions. The hypothesis

expressed by Figure 1 is that this heterogeneity of solutions in boundedly rational agents will result in some agents attending but not all of them.

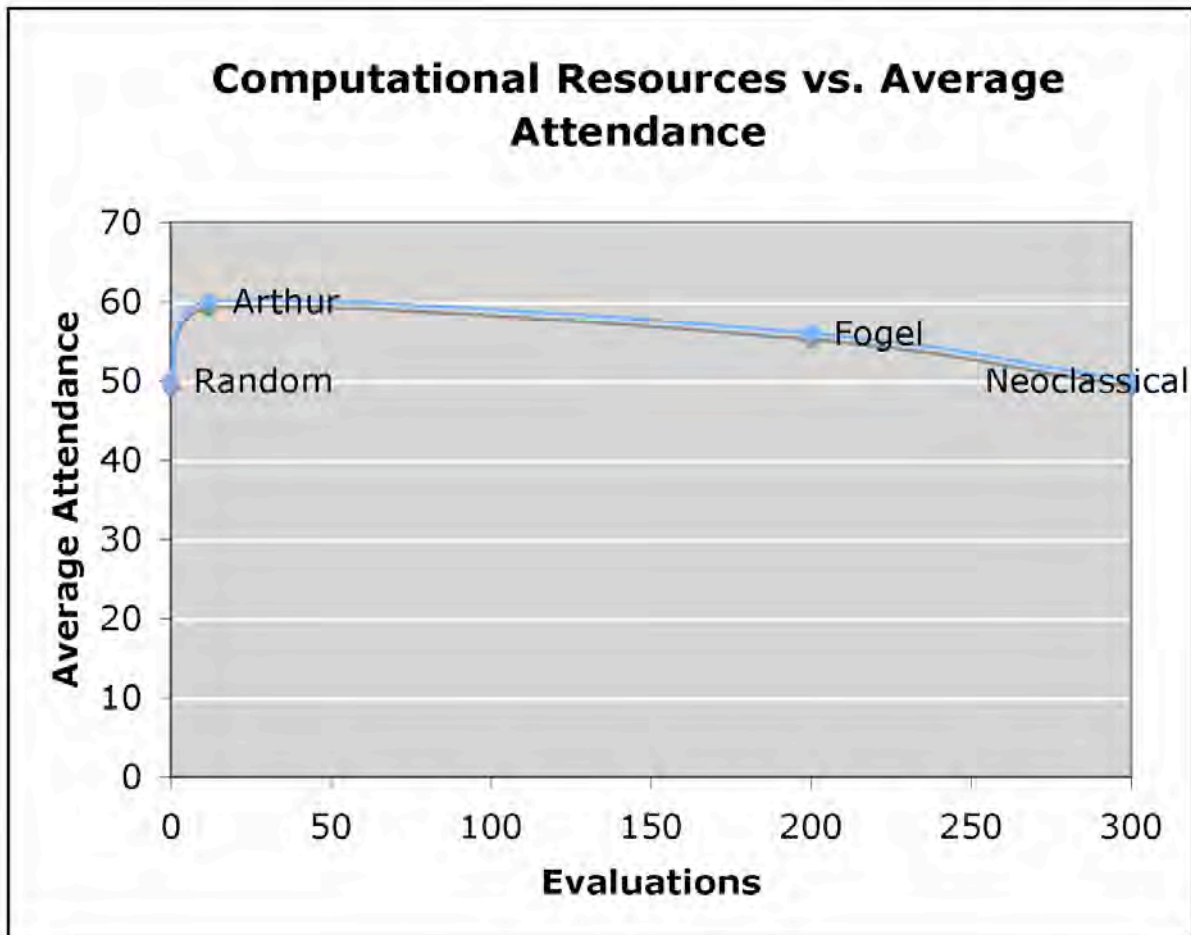


FIGURE 1 Computational Resources vs. Average Attendance.

To investigate this hypothesis we examined the El Farol Bar Problem with a group of agents where we could control the number of evaluations that each agent carried out each week. Before we get to the details of how this experiment was carried out, we will examine how we placed this model within the larger ABM-Machine Learning framework that we previously developed (Rand, 2006).

FRAMEWORK

At a high level, ABM and Machine Learning (ML) (Mitchell, 1997; Hastie et al, 2001) utilize fairly simple algorithmic structures to control their flow of operation. Roughly these algorithms can be described as: initialize the system, observe what is happening, refine the system, take actions, and repeat until time is up. As a result it is easy to examine how these systems can be integrated. Let us use the El Farol Bar Problem as an example.

Arthur's original model included a simple ML technique in it. In Arthur's model all the agents had a group of strategies. They would take this set of strategies and see which strategy would have done the best of predicting the bar attendance if they had used it in the past. Since at

each time step a new data point is generated it is possible that the actual strategy from the group of strategies that each agent will use can change at every time. This is a very simple ML technique. Initialize the population of strategies by generating some random strategies, like take last week's attendance double it, subtract the third to last week's attendance from last weeks, or take a running average of the last three weeks attendances. Then at each time step the algorithm observes how the strategies have done on the current set of training data, i.e. the previous bar attendances. After that, the algorithm refines the internal model by selecting the best strategy given the new data. Finally the algorithm acts on the strategy that reflects the refined model and repeats. As Arthur speculated (1994) and Fogel et al. showed (1999), the original El Farol Bar algorithm of a 'bag of strategies' could be replaced by another standard ML technique.

We wanted to make use of a different ML technique than the one Arthur described. There are many different ML techniques and there is no obvious best technique, but partially since it was originally suggested in Arthur's paper, we decided to investigate the use of an evolutionary algorithm, and employed the genetic algorithm (GA) as originally devised by Holland. As we have mentioned, Fogel et al. had previously explored a similar technique within the El Farol Bar Problem. The GA makes sense in this context because it has the ability to create a fairly robust time series predictor (by doing simple regression) and it is similar to Arthur's original technique, in that it considers a population of solutions, evaluates them, and decides which strategy to use. In addition the GA is often described as manipulating *schemata* and thus may be similar to the human process of induction (Holland et al., 1986) which is what Arthur's original model was intended to emulate. As a result of all of these factors we chose the GA.

We then placed the original El Farol Bar Problem and the GA within the context of the Integrated ABM-ML cycle that we had previously described (Rand, 2007). The result is illustrated in Figure 2.

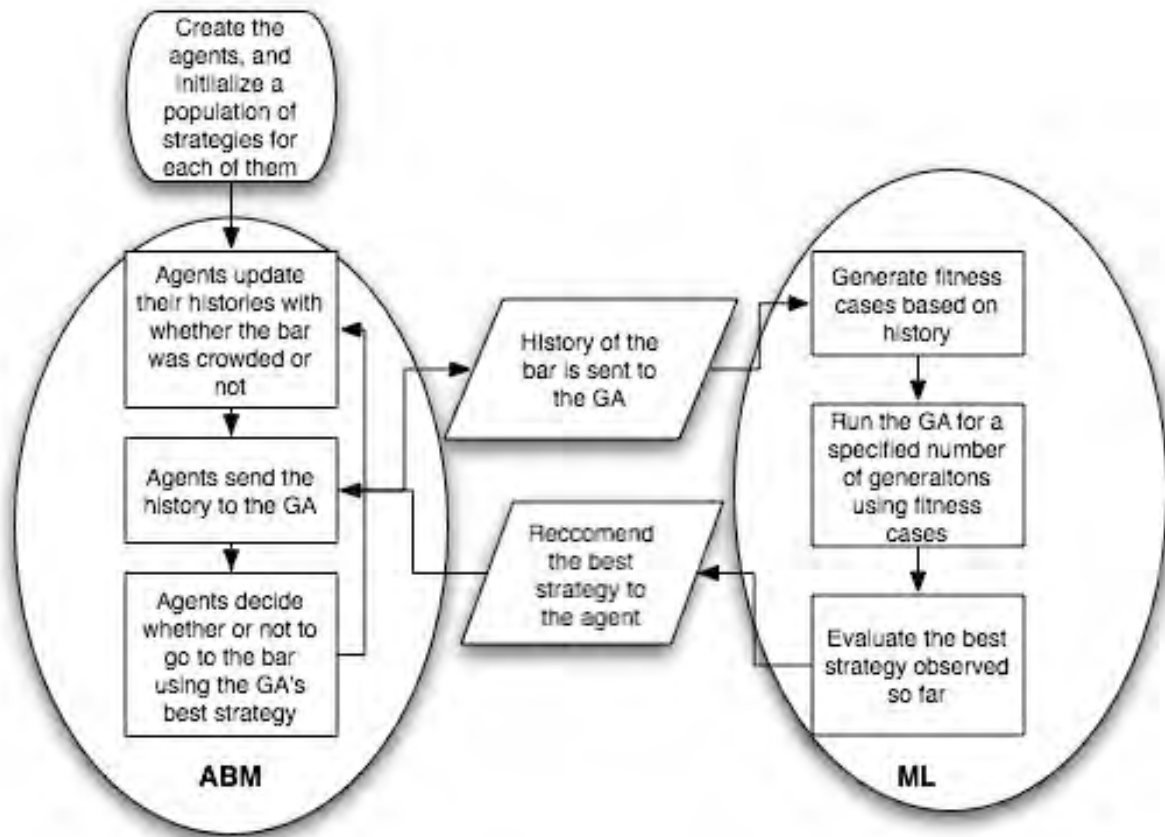


FIGURE 2 The El Farol Bar Problem and a GA within the context of the Integrated ABM-ML cycle.

EXPERIMENT

We used the framework description from Figure 2 to guide the development of an implementation of the El Farol Bar Problem in NetLogo (Wilensky, 1999). Similar to Fogel et al. (1999), we used an auto-regressive (AR) model, whereby a strategy consists of a list of real-valued numbers; these numbers are weights (in a weighted linear combination) that are used in predicting the attendance at the bar, based on attendance in previous weeks. Specifically, for a given strategy S with AR coefficients (w_0, w_1, \dots, w_L) , where L is the number of preceding weeks considered when predicting attendance, we have the following prediction formula:

$$p(S, t) = w_0 + \sum_{i=1}^L w_i a(t - i)$$

In this equation, $p(t)$ is the prediction for the attendance at week t , and $a(t - i)$ is the actual recorded attendance at week $(t - i)$. For our experiment, we fixed L at 10 weeks. Each agent had its own population of 10 strategies (initialized with weights drawn uniformly at random between -1.0 and 1.0), which we evolved over time using a real-valued genetic algorithm. Our model differs from Fogel et al. (1999) in several ways. Fogel et al. created offspring solely through asexual reproduction – each of the weights in the parents' strategy was

mutated by adding a zero mean Gaussian random variable with standard deviation 0.1. While our simulation also used this method for mutation, mutation was not the primary genetic operator. Each weight was mutated only with probability $1 / (2L)$. Instead, staying closer in form to the simple genetic algorithm (Holland, 1975), our primary genetic operator was crossover, whereby two parent strategies (lists of weights) are split and recombined to form new offspring. Fitness evaluation for a strategy consisted of measuring the sum of the prediction errors, if the strategy had been employed for the last L weeks. Each agent uses $2L$ weeks worth of attendance history, so that it can perform this evaluation. Specifically, the fitness of strategy S at current week t is defined as:

$$f(S, t) = \sum_{i=t-L}^{t-1} |p(S, i) - a(i)|$$

In this equation, we employed a linear error function where any error is weighted equally. This differed from Fogel et al., who used squared error for fitness, biasing selection toward strategies that do not make egregious errors. In order to evaluate fitness in early weeks, we provided “false” historical data – that is, $a(-1)$, $a(-2)$, ... $a(-L)$ were each initialized to random numbers between 0 and 100. We used a strictly generational GA, where a generation of 10 parents is replaced by 10 children in each generation, whereas Fogel et al.’s method evaluates 10 newly created children and 10 parents from the previous generation and chooses the 10 best strategies to create the next generation. Finally, we used tournament selection with a tournament size of 3.

The primary goal of our initial experiment was to examine the relationship between society-wide resource utilization (formalized as the bar attendance) and the amount of computational effort expended by the agents (formalized as the number of generations of strategy evolution the agents were allowed per week). Whereas Fogel et al. (1999) fixed the number of generations per week at 10, we allowed this parameter to vary from 0 to 20. Note that in the 0 generations per week case, no evolution is occurring at all. In this case, each agent has only a fixed set of 10 strategies to choose from, which is, in some respects, similar to Arthur’s model (1994). We ran the model for 500 weeks per run and carried out 30 runs for each parameter setting. Fogel et al. found that the behavior of his system reached a “steady-state” of chaotic oscillation after 100 weeks, so our choice of 500 weeks seemed sufficiently large.

In several details, our experimental set up has deviated from that of Fogel et al., and it is important to note that we are not attempting to exactly replicate their experiment or results. However, by investigating the same problem and using a similar strategy representation, our work is comparable to theirs, and differs mainly in some particular details of the evolutionary algorithm, which Fogel et al. admitted were chosen somewhat arbitrarily. Thus despite these differences, a secondary goal of our experiment was to determine if our results support the general findings of the prior work by Fogel et al (by showing them to be robust despite variations in the general evolutionary algorithm).

RESULTS AND DISCUSSION

For each run, we measured the attendance at the bar in each of the 500 weeks. Since the early weeks could be skewed by the random initial conditions, we decided to concentrate on the attendance behavior during the last 100 weeks of each run. The first metric we examined was

the mean attendance, as this was the quantity focused on by previous work (Arthur, 1994, Fogel et al., 1999). This is shown in Figure 3.

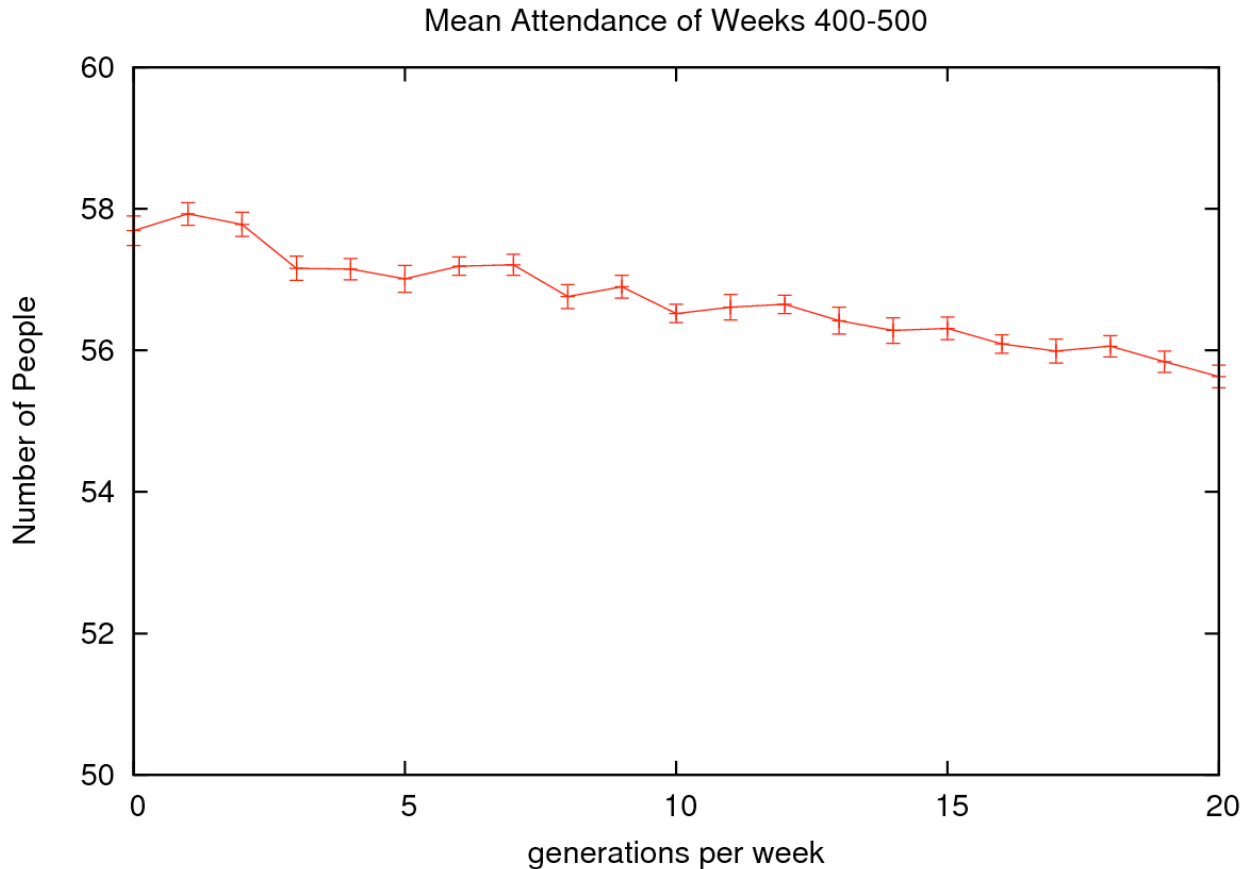


FIGURE 3 The mean (over 30 runs) attendance at the bar in the last 100 weeks versus the number of generations evolved each week. Standard error bars are shown.

Fogel et al., using 10 generations of evolution per week, found a mean attendance of 56.32. Our result for 10 generations per week was 56.52. The difference between these values is on the border of statistical insignificance, and given that we made several differing choices in our experimental setup, we feel our results support the previous findings of Fogel et al. Furthermore, our model exhibited the same non-convergent oscillations around the mean, found by Fogel et al. In addition, as we noted above, the 0 generations-per-week case is similar to Arthur's model of bounded rationality. However, as shown in Figure 3, even without evolution, our mean attendance was only 57.69, falling well short of Arthur's predicted convergence around 60.

While making comparisons to previous results is interesting, our primary goal, which was not investigated by Fogel et al. (1999) and Arthur (1994), is to determine the effect on mean attendance of varying the amount of computational effort given to agents, i.e. the number of generations-per-week. Figure 3 generally supports the hypothesis that we explained earlier. As we increase the computational power of the agent, the average attendance at the bar decreases. As can be seen in Figure 3, the difference between consecutive points on the x-axis is not always statistically significant, but the general downward trend is statistically significant. As we explained previously, one possible explanation is that as the amount of computational power

increases the agents' strategies start to converge, which drives down the average attendance at the bar. Thus Figure 3 lends some credence to this hypothesis but further work is required to confirm this hypothesis. In future work, we plan to investigate the diversity of strategies being employed across all agents.

This result was interesting, but it also prompted us to question whether mean attendance is the best metric to capture “resource utilization.” For instance, consider that there are two qualitatively different attendance ranges – if 61 people attend the bar, then no one is happy, whereas if 59 people attend the bar, then 59 people are happy. It is not clear that it makes sense to simply average the values 59 and 61, when computing resource utilization. For this reason, we defined a separate metric, which we shall call “societal benefit”, which is simply the number of happy bar patrons per week. When the bar is overcrowded, the “societal benefit” for that week is 0. We measured the average societal benefit across the last 100 weeks, and the results are shown in Figure 4.

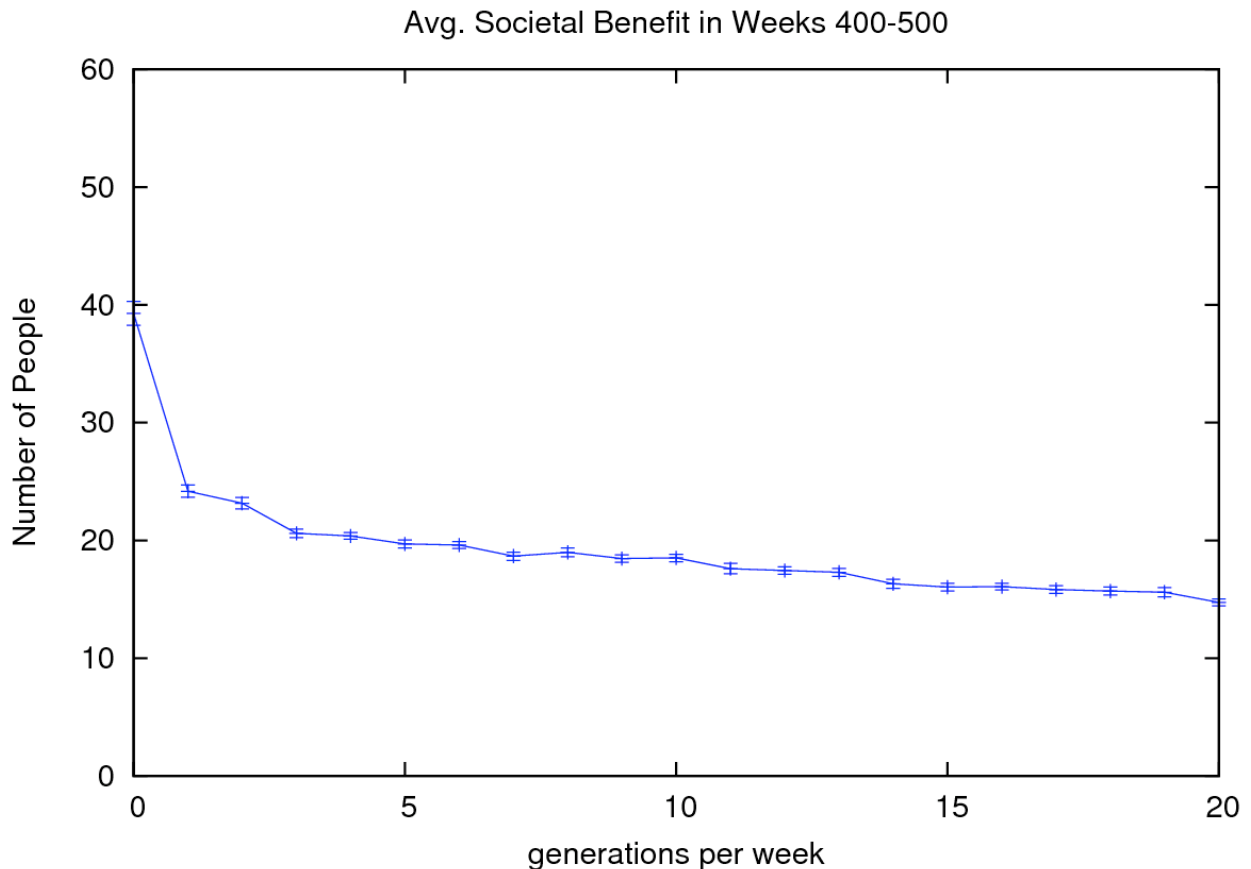


FIGURE 4 The average (over 30 runs) societal benefit in the last 100 weeks versus the number of generations evolved each week. Standard error bars are shown.

Similar to Figure 3, the relationship is that resource utilization decreases as we increase the amount of computational power given to the agents. The trend here is even more apparent. Part of this trend is obviously due to the decrease in attendance at the bar, but that is not the whole story. We know from the previous work of Fogel et al. (1999) that toward the end of the run, the average bar attendance usually fluctuated wildly. One explanation for the decrease in

societal benefit might be that these fluctuations result in fewer people attending the bar when it is below capacity. To investigate this hypothesis, we examined the standard deviation of the attendance at the bar. Figure 5 illustrates the standard deviation of the mean attendance at the bar over the last 100 weeks versus the number of generations of evolution per week. As can be seen the standard deviation does increase as the computational resources increases – that is, allowing the agents more time to evolve their strategies results in a greater amplitude of the chaotic oscillation in attendance. This supports the idea that the population is fluctuating wildly and that is why the societal benefit is decreasing. However, precisely why additional computational resources causes an increase in the size of the fluctuations remains a subject for further study.

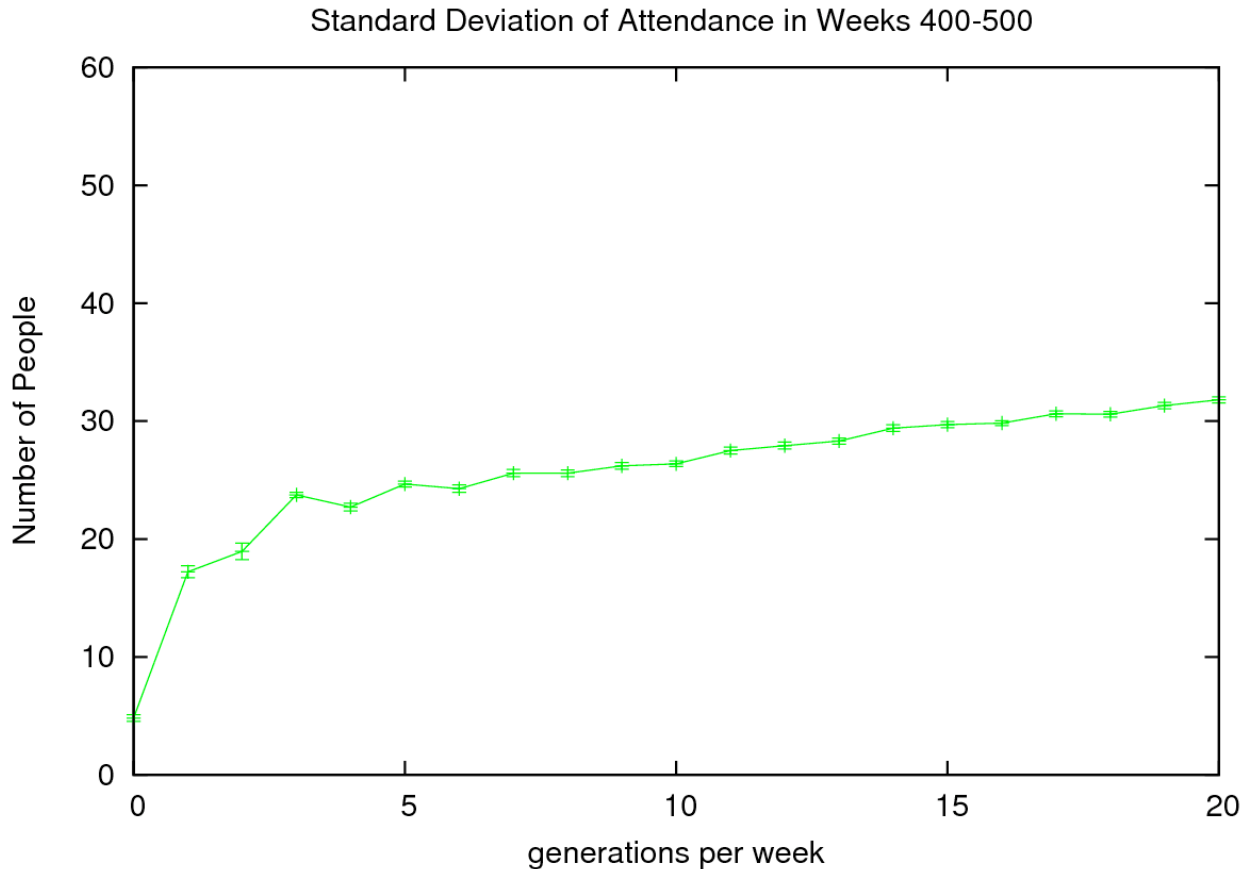


FIGURE 5 The average (over 30 runs) of the standard deviation of the mean attendance in the last 100 weeks versus the number of generations evolved each week. Standard error bars are shown.

FUTURE WORK AND CONCLUSION

These initial results are tentative, and there is more work that needs to be done to substantiate the hypotheses that we have suggested. For instance, the number of generations per week might not be the deciding factor that influences the societal benefit and the mean attendance at the bar. Another possible factor is simply the total amount of evolution. For instance, in the 5 generations per week case at week 500, each agent has undertaken 2500 evaluations, while in the 20 generations per week case at week 500, each agent has undertaken

10000 evaluations. One possible explanation is that the behavior of the 5 generations per week case at week 500 is comparable to the 20 generations per week case at week 125. Initial investigations indicate this is not the case, but additional verification is warranted. Investigating the diversity of the final strategies employed by the agents would also help substantiate some of the claims that we have discussed. Finally, it would be useful to look at the entire attendance distribution histogram rather than just aggregate measures like means and standard deviations.

Still, we have begun to investigate what the relationship between computational power and efficient resource utilization is. This experiment is a first step toward understanding if there are general claims to be made about this relationship. These tentative results indicate that it might be possible for simple machine learning algorithms to be given a limited amount of computational power and still achieve an ecology of strategies that produces a greater resource utilization than a more complex learning algorithm with greater computational power.

REFERENCES

- Arthur, W. B. (1994). "Inductive Reasoning and Bounded Rationality." The American Economic Review **84**(2): 406-411.
- Challet, D., M. Marsili and Y. Zhang (2005). Minority Games. Oxford, UK, Oxford University Press.
- Edmonds, B. (2000). "Gossip, Sexual Recombination and the El Farol bar: modeling the emergence of heterogeneity." Journal of Artificial Societies and Social Simulation. 2 (3).
- Fogel, D. B., K. Chellapilla, et al. (1999). "Inductive Reasoning and Bounded Rationality Reconsidered." IEEE Transactions on Evolutionary Computation **3**(2): 142-146.
- Hastie, T., R. Tibshirani, et al. (2001). The Elements of Statistical Learning: Data Mining, Inference, and Prediction. New York, Springer.
- Holland, J. H. (1975). Adaptation in Natural and Artificial Systems, MIT Press.
- Holland, J. H. (1995). Hidden Order: How Adaptation Builds Complexity. Reading, Massachusetts, Perseus Books.
- Holland, J. H., K. J. Holyoak, et al. (1986). Induction, MIT Press.
- Mitchell, T. M. (1997). Machine Learning, McGraw-Hill.
- Rand, W. (2006). "Machine Learning meets Agent-Based Modeling: When no to go to a Bar." In the Proceedings of Agent 2006. Chicago, IL.
- Wilensky, U. (1999). NetLogo. Evanston, IL: Center for Connected Learning and Computer-Based Modeling. <http://ccl.northwestern.edu/netlogo>.
- Wolpert, D. H., K. R. Wheeler, et al. (2000). "Collective Intelligence for Control of Distributed Dynamical Systems." Europhysics Letters. **49** (6), p.708.