

Evolutionary Algorithms and Particle Swarm Optimization for Artificial Language Evolution

Geoff Nitschke, Kobus de Bruyn, Willem van Heerden

Computational Intelligence Research Group, Department of Computer Science, University of Pretoria
Pretoria, South Africa

gnitschke@cs.up.ac.za, cjdebruyn@gmail.com, wvheerden@cs.up.ac.za

Abstract—This paper reports upon two adaptive approaches for deriving words in an artificial language simulation. The efficacy of a Particle Swarm Optimization (PSO) method versus an Artificial Evolution (AE) method was examined for the purpose of adapting communication between agents. The objective of the study was for agents to derive a common (shared) lexicon for talking about food resources in the simulation environment. In the simulation, communication was essential for agent survival and as such facilitated lexicon adaptation. Results indicated that PSO was effective at adapting agents to quickly converge to a common lexicon, where, on average, one word for each food type was derived. AE required more method iterations to converge to a common lexicon that contained, on average, multiple words for each food type. However, there was greater word diversity in the lexicon converged upon by AE evolved agents, compared to that converged upon by PSO adapted agents.

Index Terms—Artificial Language, Particle Swarm Optimization, Evolutionary Algorithm, Artificial Life.

I. INTRODUCTION

Artificial Life research has often applied biological principles and the methodology of building artificial systems to understand the origins and evolution of communication [17], [10], [3]. In such research, language is viewed as a complex adaptive system which emerges, and self-organizes in a bottom-up means from the local interactions between agents. Research in *artificial language evolution* thus applies adaptive processes in order to explore *language as it could be*, in much the same way as artificial life explores *life as it could be* [9]. The use of a synthetic methodology has been applied to model the evolution of communication in computer simulation [16], [11], [21], [15], as well as using situated and embodied robotic agents operating in real world environments [1], [18], [20], [19]. Synthetic methodologies used to study the origins and evolution of language include *Evolutionary Algorithms* (EAs) [6], *Artificial Neural Networks* (ANNs) [14], and *rule-based agents* [15]. Such bottom-up synthetic simulations are indispensable since they allow researchers to readily study language evolution, and formulate and test hypotheses. Without bottom-up simulations, the study of the emergence of language has been problematic since language is a complex, nonlinear, and analytically intractable system.

There are numerous examples of synthetic bottom-up agent based simulations that investigate language evolution. In the *Talking Heads* simulation, a common (shared) lexicon was derived between robots with rule-based controllers. The robots played a series of *naming games* [16] in simulated and

physical environments [18]. An iterative process facilitated the emergence of a common lexicon with positive feedback loops, where words used by multiple robots were reinforced, and propagated through successive generations. Similarly, in the research of de Boer [4], rule-based agents interacted in the context of iterative imitation games in order to derive a common lexicon. Results supported related work [16], and indicated that a common lexicon of vowels emerged in an agent group as a consequence of local interactions between agents. Cangelosi and Parisi [2] simulated a population of agents with evolving ANN controllers, for the purpose of deriving a common lexicon that aided agent survival. The objective was for agents to evolve signals to help other agents identify edible versus poisonous mushrooms in the environment. Results demonstrated that agent communication evolved as a product ancillary to the need for agents to evolve categorizations for mushrooms.

The application of *Particle Swarm Optimization* (PSO) to adapt agents within artificial language evolution simulations has, however, received relatively little research attention. Furthermore, there has been a lack of research that compares PSO and EA methods for adapting agent behaviors, where agents must derive a common lexicon in order to communicate.

This paper's research applied PSO and EA methods for agent adaptation in a *talking game* task. This task required that agents communicate (adaptively derive a common lexicon) in order to consume food, survive, and procreate in an artificial environment. In this study *common lexicon* refers to one word (or a set of *similar words*¹) used by at least two agents to describe a given food type. Section II describes these elements of the talking game task and the talking game process.

PSO and EA were selected as the comparative methods since both approaches have been successfully applied as adaptation methods in agent-based simulations [7].

To the best of the authors' knowledge, there has not yet been any research that compares the efficacy of a PSO versus an EA method for agent adaptation in artificial language evolution simulations. This paper addresses the general research goal of ascertaining which adaptive methods, implemented in the context of an agent-based simulation, are most appropriate for studying the origins and emergence of communication.

Using artificial language simulations to study the emergence

¹In this study, a set of *similar words* for a given food type (word similarity is defined in section II) is analogous to one word for the food given type.

of communication is important not only to linguistics, but for numerous engineering disciplines. That is, understanding the fundamental mechanisms that lead to new forms of communication, may allow such mechanisms to be applied to the behavioral design of agents (simulated) or robots (physical) that must interact with each other in order to accomplish a task. This notion is supported by the successful development and application of various forms of biologically inspired computation such as neural networks, genetic algorithms and swarm intelligence algorithms [23].

A. Research Goal and Hypotheses

- *Research Goal:* To conduct a comparative study that evaluates a PSO versus an EA method for adapting agents such that they derive a common lexicon.
- *Hypothesis 1:* The PSO, comparative to the EA method, will yield a statistically significant higher average fitness². This hypothesis is supported by related research [5], [22], [12] reporting that PSO, comparative to EAs, often facilitates (for small population sizes), convergence to optimal regions of the search space.
- *Hypothesis 2:* The PSO method is appropriate for adapting agents to converge to a common lexicon with, on average, one word for each food type. Whereas, the EA will evolve agents that converge to a common lexicon containing multiple words for each food type. This hypothesis is supported by related research that found that PSO, compared to EAs, often converges upon less diverse (genetic and behavioral) solutions [5], [22], [12].

II. TASK: THE TALKING GAME

The *talking game* task places N agents, and X *red*, Y *green*, and Z *blue* food units, at random locations on a two dimensional grid of 10×10 cells. Each agent's *genotype* is a character set which represents the agent's word *word*³ for one food type (food types are: *red*, *green*, and *blue*). *Common lexicon* refers to the case where at least two agents share *one*, or a set of *similar words*, for one of the food types. Hence, an example of a common lexicon is if three sub-groups n_0 , n_1 , and n_2 (where, n_0 , n_1 , and n_2 constitute the N agents of the group), each share one word for *red*, *green*, and *blue* food types, respectively.

Measuring word similarity is discussed in section II-A. A PSO or EA method was used to adapt agents to accomplish the talking game task. This task was accomplished when all agents converged upon a common lexicon.

In each simulation, agents began with zero fitness, no preferred food type, and a word for its preferred food type (initially a random character set). *Preferred food type* refers to a food type that an agent will attempt to talk about (and thus consume). When an agent first encounters a food unit, the food's type (*red*, *green*, or *blue*) becomes its preferred food

type. Thereafter, when an agent moves to a cell adjacent to its preferred food type, it will try to begin a talking game. In the talking game, an agent's *genotype* and its *word* for its preferred food type are identical. There cannot be more than one food unit at a given location, and agents cannot occupy the same location as a food unit. Agents encounter food units when they move to a grid cell adjacent to a food unit cell.

A. Talking Game Procedure

- 1) *Agents move:* All agents move concurrently in a random direction to adjacent grid cells (*north*, *south*, *east*, *west*, *north-east*, *north-west*, *south-east*, or *south-west*). Conflicts between m agents moving to the same cell are resolved by $m-1$ agents randomly moving to other cells, and doing so until all conflicts are resolved.
- 2) If two agents occupy grid cells adjacent to a food unit cell, then a talking game will start. These two agents become *talking agents*. Go to step 4.
- 3) If more than two agents occupy grid cells adjacent to a food unit cell, then the two agents with the *most similar* genotypes (step 3(a)) start a talking game (step 4). These two *most similar* agents become *talking agents*.
 - a) *Genotype Similarity:* An *edit distance* metric [13] measures the similarity between two agent genotypes a and b . Similarity is the number of genes (characters) a and b have in common divided by genotype length (all genotypes have equal length). If the similarity between a and b is greater than a given *similarity threshold*, a and b are considered similar. For example, given a 0.7 similarity threshold, the similarity between $a = \textit{mouse}$, and $b = \textit{house}$ is 0.8, since 4 out of 5 characters are the same. Hence, a and b are considered similar. A similarity measure of 1.0 indicated that two genotypes were most similar, and 0.0 indicated that two genotypes were most dissimilar.
- 4) *Talking Game:* One agent is randomly selected as the *speaker*, and the other agent is the *listener*.
 - a) If the speaker's *preferred food type* is the same as the food type in the food cell adjacent to the talking agents, then go to step 4(e).
 - b) If a speaker does not yet have a preferred food type yet, then the type of the food unit to which the talking agents are adjacent becomes the speaker's preferred food type. Go to step 4(e).
 - c) If the speaker's preferred food type is different from the food type in the food cell adjacent to the talking agents, the agents switch speaker and listener roles. If speaker and listener roles have already been switched, then go to step 4(i).

²*Fitness* and *energy* are the same, and the terms are used interchangeably.

³*Genotype* and *word* are the same and used interchangeably.

- d) If the speaker's (previously the listener) preferred food type is the food type in the food cell adjacent to the talking agents, then go to step 4(e).
 - e) *Speaker talks*: The word spoken to listener is the speaker's word for its preferred food type.
 - f) *Listener guesses*: The listener guesses to which food type the speaker's word refers. If the speaker and listener genotypes are *similar* (step 3(a)), then the listener will guess its own preferred food type. Otherwise, the listener will guess a random food type. For example, assume the talking agents are adjacent to food cell containing a red food type. The word of speaker *a* for its preferred food type (red) is *aka*. The word of listener *b* for its preferred food type (green) is *aki*. Assume that *aka* and *aki* are calculated as *similar*. Listener *b* will guess food type *green*. In this case, agents *a* and *b* had similar genotypes, however, they could not agree on the food type in the cell, and would not be rewarded. This is an example of two agents adopting similar words for different food types.
 - g) If the listener's guess is the same as the speaker's preferred food type, both agents consume the food and receive a fitness reward equal to one. The food unit is then removed from the grid.
 - h) If listener's guess is not the speaker's preferred food type, and speaker and listener roles have not been reversed, then agents switch speaker and listener roles, and the talking game is restarted. Go to step 4(d).
 - i) Talking game ends.
- 5) If the agent group's (N agents) lifetime (a total of 200 talking games) has transpired, then apply either the EA (section III-B) or PSO method (section III-A) to adapt agents. If it is the end of an agent group's lifetime, a new set of food units and agents are initialized in random locations in the environment. Otherwise, go to step 1.

Hence, over the course of one EA generation or PSO method iteration (an agent's lifetime), agents adopt preferred food types (red, green, or blue), and words to label these preferred food types. Agents with the most similar genotypes will be those selected to talk (assuming more than two agents are adjacent to a food cell at a given iteration). Given that a speaker has the same preferred food type as the food cell to which it is adjacent, then the listener (with the most similar genotype) will guess its preferred food type.

However, if the listener does not have a preferred food type it will select and guess the food type in the food cell to which it is adjacent. This facilitates the adoption of similar words (or one word) for a given food type. If agents with similar genotypes cannot talk (and consume food) due to both having different preferred food types, then these agents will,

on average, be less fit at the end of their lifetimes. As a result, such agents will have less chance of being selected at the end of an EA generation (or PSO iteration) and having their genotypes propagated.

That is, agents with similar genotypes suggests that these agents are successful at consuming food, which is a result of the agents with similar genotypes having the same preferred food type. The genotypes of the fittest agents (with the same preferred food types) will then be adapted by the EA and PSO methods and be propagated such that agents converge to similar words (or the same word) for a given food type. The highest task performance is attained when all agents converge to the same word for a given food type.

B. Task Performance Evaluation

- 1) *Average Fitness*: An agent was rewarded one unit of fitness for each food unit consumed. Each agent's fitness was recorded at the end of each simulation (n EA generations, or PSO iterations), and an average fitness was calculated for N agents in the group.
- 2) *Size of Common Lexicon*: This is the sum of the average number of words derived by the agents (adapted with the EA or PSO method) for each food type. A given food unit can only be consumed if two agents agree upon the *food type* of the food unit that is to be consumed. Agents maximized their food consumption if all agents used a common lexicon (containing one word, or multiple *similar* words for each food type). The case where all agents use the same word for a given food type is an example of the most effective (smallest) common lexicon. When all agents use a different word for a given food type, this is an example of the most ineffective (largest and non-shared) lexicon. In this case, agents would be unable to communicate and consume food.

The size of a common lexicon is the number of *dissimilar* of words derived by agents for each food type. That is, the effectiveness of convergence upon a common lexicon was simply indicated by the number of words derived for each food type. If agents derived u_0, u_1 and u_2 words for red, green, and blue food types, respectively, and the words within the sets u_0, u_1 and u_2 , were calculated as being *similar*, then the agents were considered to have derived one word for each food type. This is an example of the most effective convergence to a common lexicon. However, consider that the agents derived u_0, u_1 and u_2 words for the red, green, and blue food types, respectively, and at least two words in the word sets u_0, u_1 and u_2 , were calculated as being *dissimilar*. In this case agents were considered to have been less effective at converging upon a common lexicon. Section II-A describes methods to measure word similarity.

III. METHODS

The efficacy of an EA [6] and PSO [7] method for adapting agents was tested in the talking game task (section II).

A. PSO: Particle Swarm Optimization

PSO methods model a set of potential solutions as a swarm of particles that move about in a virtual search space. Each particle has a *position* and a *velocity* vector that is updated at each algorithm iteration [7]. A velocity update consists of:

- 1) A *cognitive* component (c_1), which uses each particle's personal best (*pbest*) solution.
- 2) A *social* component (c_2), which uses a neighborhood best solution (*nbest*).
- 3) An *inertia* coefficient (ϕ), which slows particle velocity over time to facilitate swarm convergence.

The influence of the c_1 , c_2 , and ϕ components on the adaptive process of many PSO methods, including the PSO method used in this study, is discussed in Engelbrecht [7].

PSO Type: This PSO method uses a global best (*gbest*) neighborhood structure for the social component of the velocity update [7]. Hence, the entire swarm constitutes the neighborhood of the particles. The social information is the global best, that is, the current best position found by the entire swarm. At each iteration of the PSO method, a particle's velocity is added to its current position, and particles are evaluated using the talking game (section II).

Particles: The PSO method implements one particle for each agent in the simulation. The position vector of each particle represents each agent's word for its preferred food type. Each position vector element is a character in a five character set (table I). Each element in the position vector is randomly initialized within the range: a to z . This character range is converted to ASCII values (97 to 122). Since the PSO method operates using floating point numbers, the adapted elements in a position vector were rounded to the nearest integer so that distinct characters could be represented as agent words. The inertia weight w was implemented to linearly decrease from 0.3 to 0.01, as a function of the PSO method iterations. A linearly decreasing w was selected based on the success of such an approach in related research [7].

Particle Evaluation: At each PSO method iteration, 200 talking games (section II-A) were played. These games evaluated agent words (particles), assigning each a fitness. The assignment of fitness to all agents constituted one iteration in the PSO method. All particle positions were updated at the end of each PSO iteration.

B. EA: Evolutionary Algorithm

The EA method used was adapted from the evolutionary algorithm described by Eiben and Smith [6].

Genotypes: The N agents were represented by N genotypes in the EA population. Each genotype was encoded as a string of n integer values representing the characters in each agent's word for its preferred food type.

Selection and Recombination: Tournament selection [6] was used, such that the fittest genotype was selected from a tournament size of 0.3 (that is, 30% of genotypes). Parent genotypes were randomly selected, where each pair produced one child genotype, until there were enough child genotypes to completely replace the previous generation (parent population). *One-point crossover* [6] was used to recombine the genotypes of agent pairs at each generation.

Mutation: After recombination, mutation was applied with a given probability to each character in a given genotype. A mutated character was changed to another random character in the range [a, z] (that is, ASCII values in the range: [97, 122]). The mutation rate (applied per gene), linearly decreased from 0.09 to 0.001, as a function of generations. This approach was selected given the success of linearly decreasing mutation reported upon in related research [7].

IV. EXPERIMENTS

Experiments applied the PSO and EA methods for adapting agent behaviors. Each experiment executed the PSO or EA method for 1000 iterations (generations). Each method iteration consisted of 200 talking games. These 200 talking games represented an agent lifetime. At each agent lifetime iteration, each agent either moved about the grid, or engaged in a talking game with another agent. At the end of each lifetime, each agent's performance was evaluated and a fitness assigned. An averages fitness and number of words derived for each food type was calculated for all agents, over all PSO and EA method iterations and for 30 (PSO and EA) simulation runs.

Table I presents the simulation, EA, and PSO parameter settings used to attain the results presented in section V. Experiments used to derive these parameter values used the same experimental setup as PSO and EA method comparison experiments (section IV-C), except that the affect of varying individual parameter values was observed (sections IV-A and IV-B). Statistical significance of difference between parameter tuning results were calculated using an independent t test [8].

A. Parameter Tuning in the PSO Method

For the PSO method's *cognitive* (c_1) and *social* (c_2) terms, and the *inertia weight* (ϕ), 10 values were tested in increments of 0.1 in the range [0.1, 1.0].

Cognitive Term (c_1): Exploratory experiments indicated that setting c_1 to 0.1 resulted in agents with a high average fitness. Increasing c_1 towards 1.0 resulted in agents yielding a comparatively low average fitness. Setting c_1 to 1.0 resulted in poor convergence upon a common lexicon (a high average number of words was derived for each food type). Decreasing the value of c_1 towards 0.1, adapted agents such that they converged to a common lexicon with, on average, few words for each food type.

Social Term (c_2): Setting c_2 to 1.0, resulted in agents yielding a higher average fitness comparative to when c_2 was set to lower values. Setting c_2 to 1.0 resulted in agents deriving a common lexicon containing few words for each

food type. As the c_2 value was decreased towards 0.1, less convergence towards a common lexicon was observed.

Inertia Weight (ϕ): Exploratory experiments indicated that there was no statistically significant difference between the average fitness of agents adapted using each of the ϕ values tested. Hence a ϕ value that linearly decreases as a function of PSO method iterations was selected. The ϕ value decreased from 1.0 (by a value of 0.1) every 100 PSO iterations.

In parameter tuning experiments, c_1 and c_2 values equal to 0.5, together with a linearly decreasing ϕ value, were found to adapt agents that yielded the highest fitness.

B. Parameter Tuning in the EA Method

For the EA method parameters, *tournament size* and *mutation rate*, 10 values were tested in 0.1 increments in the range [0.1, 1.0]. Exploratory experiments also tested the EA method without crossover and without mutation.

Tournament Size: Results indicated that for all tournament sizes tested, there was no statistically significant difference in the evolved agents' average fitness, and number of words derived for each food type. A tournament size of 0.3 was selected since this value provided (for all numbers of agents and food units tested) a good balance between the EA's exploration versus exploitation of the search space.

Mutation Rates: Results indicated that, for all mutation rates tested, there was no statistically significant difference in the evolved agents' average fitness and number of words derived for each food type. Based on related research [7], a mutation rate that linearly decreased as a function of generations was used. The mutation rate decreased from 0.09 to 0.001, by a value of 0.001 every 100 generations.

Mutation and No Crossover: Results of testing the EA with mutation and without crossover indicated that evolved agents yielded a relatively high average fitness (comparative to the EA method using crossover and no mutation). However, in this case, evolved agents were less effective at converging upon a common lexicon. That is, EA evolved agents converged to an average of: [10.06 (2.43), 8.37 (1.92), 6.57 (2.01)] words for red, green, and blue food types, respectively. Values given in parentheses are standard deviations.

Crossover and No Mutation: Results of testing the EA method with crossover and without mutation indicated that evolved agents yielded a low average fitness (comparative to the EA method using mutation and no crossover).

The mechanisms causing the EA method, with mutation only, to evolve agents with a high average fitness, and to derive a common lexicon, containing on average one word for each food type, is the subject of ongoing research. Similarly, the result of the EA, with crossover only, evolving agents with a high average fitness that also derive a common lexicon, is also the subject of current research. These results are not examined

TABLE I
PSO, EA, AND SIMULATION PARAMETERS.

| PSO Parameters | |
|--|--|
| Swarm Size | 25 |
| Number of PSO iterations | 1000 |
| Inertia Weight Operator | Linearly decreasing |
| Inertia Weight (ϕ) | [0.01, 0.30] |
| Cognitive Term (c_1) | 0.5 |
| Social Term (c_2) | 0.5 |
| Velocity Max ($VMax$)/Initialization | 4 / 0 |
| Neighborhood topology | <i>gbest</i> |
| Neighborhood size | Swarm |
| EA Parameters | |
| Population size | 25 |
| Number of EA generations | 1000 |
| Mutation operator | Linearly decreasing |
| Mutation rate per gene (σ) | [0.001, 0.09] |
| Selection operator | Tournament |
| Tournament Size | 0.3 |
| Gene value initialization | [97, 122] |
| Recombination operator | One-point crossover |
| Talking Game Simulation Parameters | |
| Iterations per talking game | 1 |
| Agent lifetime | 200 talking games |
| Talking games per simulation | 200000 |
| Simulation runs per experiment | 30 |
| Environment width/length | 10 x 10 |
| Initial agent positions | Random |
| Initial agent fitness | 0 |
| Edit distance similarity threshold | 0.7 |
| Genotype (Word) character set | [a, z] |
| Genotype (Word) size | 5 Characters |
| Food Distribution | $\frac{1}{3}$ Red, $\frac{1}{3}$ Green, $\frac{1}{3}$ Blue |
| Fitness yield per food unit | 1 |
| Number of agents/food units | 25 / 30 |

here, since the impact of mutation and crossover in the given EA are not the focus of this study.

C. Task Performance Comparisons

Task performance is the *average fitness* and *average number of words* derived for each food type. These averages were calculated, for all agents, at the end of a PSO or EA process (1000 iterations or generations), and over 30 simulation runs. PSO and EA method task performance was compared for varying numbers of agents and food units (tested at increment values of 5, in the range [5, 50]). For all experiments, methods were compared for a distribution of equal portions of the red, green, and blue food types (table I).

Experimental Objective: To ascertain which method maximizes average fitness, and minimizes the average number of words that are derived for each food type, for all numbers of agents and food units tested.

- *Average Fitness*: This task performance measure tested hypothesis 1 (section I-A). That is, whether the PSO, comparative to the EA method, will yield a statistically significant higher average fitness.
- *Lexicon Convergence*: The average number of words derived for each food type, was used to test hypothesis 2 (section I-A). It is hypothesized that the PSO, versus the EA method, is more appropriate for adapting agents that converge to a common lexicon with few words for each food type, and the EA will adapt agents with comparatively poor converge. That is, EA evolved agents will be less effective at converging to a common lexicon, on average deriving, more words per food type. The word similarity measure is described in section II-A.

V. RESULTS

This section presents task performance results for agents adapted by the EA and PSO methods. That is, *average fitness* and *average number of words* for each food type (common lexicon convergence) results (sections V-A and V-B, respectively) for a range of agent group sizes and food unit numbers.

To gauge task performance results of the comparative methods, statistical tests were applied to the PSO and EA result data. The Kolmogorov-Smirnov test [8] confirmed that the PSO and EA result data conformed to normal distributions. In order to determine if there was a statistically significant difference between the task performances of PSO and EA adapted agents, an independent t test [8] was applied. A statistical significance of 0.05 was selected, and the null hypothesis stated as the data sets not significantly differing.

A. Average Agent Fitness

Figure 1 (left and right) presents the average fitness for the PSO and EA methods, graphed for varying numbers of food units and agent group sizes, respectively. These average fitness results are graphed for increments of five in numbers of food units and group sizes, in the range [5, 50].

Statistical tests indicated that, for all numbers of food units tested, that PSO adapted agents yielded a statistically significant higher average fitness, comparative to that yielded by EA adapted agents, except for environments with five food units (figure 1, left). This statistical significance of difference in average fitness also held for all agent group sizes tested.

These results indicate that the PSO method, comparative to the EA, was effective at adapting agents such that they converged to a common lexicon containing, on average, one word for each food type. This enabled PSO adapted agents to successfully participate in a greater number of talking games, and thus consume more food and receive more fitness. These results support hypothesis 1 (section I-A). That is, the PSO, compared to the EA method, is more appropriate for adapting agents that yield a higher average fitness (for all numbers of food units and agents tested).

Contributing to this result, was that the PSO method converged to fit solutions (as a consequence of convergence to a common lexicon containing one word for each food type)

with greater rapidity, comparative to the EA method. This is supported by related research [12], [5], [22], that demonstrates that PSO often has the advantage of converging, faster than EAs, to optimal regions of the search space. In this study, quick convergence to a common lexicon (section V-B), was beneficial for PSO adapted agents, in that the few words derived for each food type allowed many agents to successfully participate in talking games, consume food, and increase their fitness. This was not the case in EA adapted agents, since such agents could not participate in talking games with the same degree of success (given the greater number of words used for each food type), and thus could not consume as much food and increase their fitness by as much. Results for the average number of words derived for each food type by PSO adapted agents (section V-B), supports this statement.

B. Lexicon Convergence

Figures 2, 3, and 4 present, for the EA and PSO methods, the average number of words derived for each food type. In figures 2, and 4 (right), the average number of words, for red, green, and blue food types, respectively, is given for increments of five in agent group sizes in the range [5, 50]. Similarly, for figures 3, and 4 (left), the average number of words derived for red, green, and blue food types, respectively, is given for increments of five food units in the range [5, 50]. In figures 2 to 4, words used to calculate the average number of words graphed were defined as being *dissimilar* according to the word similarity measure defined in section II-A.

A statistical comparison of the average number of words derived, by PSO adapted versus EA evolved agents indicated that PSO adapted agents converged upon a common lexicon containing a statistically significant lower average number of words, comparative to that converged upon by EA evolved agents. This statistical significance of difference between the average number of words derived held true for all numbers of agents and food units tested. Exceptions were observed in environments containing 5 agents for the red (figure 2, left) and blue food types (figure 4, right).

C. Results Discussion

These results indicate that, for all food types, the average number of words derived by EA evolved and PSO adapted agent groups increased with agent group size. Furthermore, for all food types, PSO adapted agents were effective and efficient at converging upon a common lexicon. That is, for all agent group sizes and numbers of food units tested, within 400 method iterations, PSO adapted agents converged to an average of 1.20, 1.10, and 1.0 words for *red*, *green* and *blue* food types, respectively. This result was observed for all PSO method simulations. In contrast, EA evolved agents were not as effective or efficient at converging upon a common lexicon. Rather, EA evolved agents derived a common lexicon containing a higher average number of words per food type for all numbers of agents (figures 2, and 4, right) and food units (figures 3, and 4, left) tested. The exception was for the average number of words derived for red (figure 2, left) and blue (figure 4, right) food types for groups of 5 agents.




Fig. 1. EA/PSO: Average fitness calculated for varying numbers of food units and 25 agents (left), varying agent group sizes and 30 food units (right).




Fig. 2. EA/PSO: Average number of words derived for *red* (left) and *green* (right) food versus agent group size and 30 food units.




Fig. 3. EA/PSO: Average number of words derived for *red* (left) and *green* (right) food versus number of food units and 25 agents.




Fig. 4. EA/PSO: Average number of words for *blue* food versus number of food units and 25 agents (left), agent group size and 30 food units (right).

Furthermore, for each food type, PSO adapted agents derived words with greater similarity, comparative to those derived by EA adapted agents. That is, over all PSO adapted processes (30 simulation runs), average word similarity (section II-A) for *red*, *green*, and *blue* food types was calculated as: 0.10, 0.03, and 0.0, respectively. Comparatively, over all EA adaptation processes (30 simulation runs), average word similarity for *red*, *green*, and *blue* food types was calculated as: 0.37, 0.97, and 0.27, respectively. These average word similarity results, and the statistically significant higher average fitness of PSO adapted agents, comparative to EA evolved agents supports the hypothesis that PSO is more appropriate for facilitating convergence of an agent group to a common lexicon (hypothesis 2 in section I-A). That is, having (on average) one word per food type in the common lexicon, allowed more PSO adapted agents to successfully engage in talking games and thus consume more food and increase average fitness. A detailed analysis of the underlying mechanisms responsible for the statistically significant higher average fitness and convergence to a common lexicon by PSO adapted agents is currently work in progress.

These lexicon convergence results are also supported by related research that has demonstrated that PSO is often more effective, comparative to EAs, for rapidly converging to beneficial regions of search space. This has especially been the case for PSO and EA method comparisons that use small population sizes [5], [22], [12]. In this study, the PSO and EA methods used 25 particles or genotypes (table I).

VI. CONCLUSIONS

This research compared the efficacy of a PSO versus an EA method for deriving and adapting communication in a simulation where agents played talking games. A talking game required that agents communicate in order to consume red, green, and blue food types and thus survive in the simulation environment. The task performances of the PSO and EA methods were comparatively evaluated according to the criteria of average fitness and convergence upon a common lexicon. Average fitness was equated with the average amount of food consumed over the course of a PSO or EA adaptive process. Lexicon convergence was measured by the average number of words derived by the agents for each food type.

Results indicated that PSO adapted agents, for all numbers of agents and food units (environments) tested, yielded a higher statistically significant fitness comparative to the EA adapted agents. Furthermore, PSO adapted agents were effective and efficient (taking, on average, less than 400 method iterations) at converging to a common lexicon (comprising, on average, one word for each food type). Whereas, the EA method evolved agents that were less efficient and effective at converging upon a common lexicon.

Future research will investigate the impact of other adaptive methods as well as larger agent groups upon average fitness and convergence to a common lexicon. Also, the impact of more complex task and environment constraints will be investigated. For example, agent survival, and emergent social phenomena such as cooperation and competition, potentially

depend upon agents deriving different words for actions suited to different situations. Deriving words that mean *help*, *go away*, and *lets eat* may facilitate the emergence of cooperation within linguistic similar agent groups and competition between linguistic dissimilar agent groups. This will contribute to the general research goal of ascertaining the types of adaptive methods (and simulation environments) that are most appropriate for investigating the origins and emergence of natural and artificial forms of communication.

REFERENCES

- [1] A. Billard and K. Dautenhahn. Experiments in learning by imitation: Grounding and use of communication in robotic agents. *Adaptive Behavior*, 7(3):411–434, 1999.
- [2] A. Cangelosi and D. Parisi. The emergence of a language in an evolving population of neural networks. *Connection Science*, 10(1):83–97, 1998.
- [3] A. Cangelosi and D. Parisi. Computer simulation: A new scientific approach to the study of language evolution. In *Simulating the Evolution of Language*, pages 3–21. Springer, Berlin, Germany, 2002.
- [4] B. de Boer. Generating vowel systems in a population of agents. In *Proceedings of the European conference on Artificial Life*, pages 503–510, Brighton, 1997. MIT Press.
- [5] R. Eberhart and Y. Shi. Comparison between genetic algorithms and particle swarm optimization. In *Evolutionary Programming VII*, pages 611–616. Springer, Berlin, Germany, 1998.
- [6] A. Eiben and J. Smith. *Introduction to Evolutionary Computing*. Springer, Berlin, Germany, 2003.
- [7] A. Engelbrecht. *Fundamentals of Computational Swarm Intelligence*. Wiley, New York, USA, 2006.
- [8] B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes*. Cambridge University Press.
- [9] C. Langton. *Artificial Life: An Overview*. MIT Press, Cambridge, USA, 1995.
- [10] M. Nowak, N. Komarova, and P. Niyogi. Computational and evolutionary aspects of language. *Nature*, 417(1):611–617, 2002.
- [11] D. Parisi. An artificial life approach to language. *Brain and Language*, 59(1):121–146, 1997.
- [12] J. Pugh and A. Martinoli. Parallel learning in heterogeneous multi-robot swarms. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 3839–3846, Singapore, 2007. IEEE Press.
- [13] E. Ristad and P. Yianilos. Learning string-edit distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(1):522–532, 2009.
- [14] D. Rumelhart, G. Hinton, and R. Williams. *Learning internal representations by error propagation: Parallel Distributed Processing*. MIT Press, Cambridge, 1986.
- [15] K. Smith, H. Brighton, and S. Kirby. Complex systems in language evolution: the cultural emergence of compositional structure. *Advances in Complex Systems*, 6(4):537–558, 2003.
- [16] L. Steels. A self-organizing spatial vocabulary. *Artificial Life*, 2(2):319–332, 1996.
- [17] L. Steels. The synthetic modeling of language origins. *Evolution of Communication*, 1(1):1–34, 1997.
- [18] L. Steels. The origins of syntax in visually grounded robotic agents. *Artificial Intelligence*, 103(1-2):133–156, 1998.
- [19] L. Steels. Evolving grounded communication for robots. *Cognitive Sciences*, 7(7):308–312, 2003.
- [20] L. Steels and F. Kaplan. AIBO's first words: The social learning of language and meaning. *Evolution of Communication*, 4(1):3–32, 2001.
- [21] J. Suzukia and K. Kaneko. Imitation games. *Physica D: Nonlinear Phenomena*, 75(1):328–342, 2002.
- [22] J. Vesterstrom and R. Thomsen. A comparative study of differential evolution, particle swarm optimization, and evolutionary algorithms on numerical benchmark problems. In *Proceedings of the Congress on Evolutionary Computation*, pages 1980–1987, Portland, USA, 2004. IEEE Press.
- [23] K. Wagner, J. Reggia, J. Uriageraka, and G. Wilkinson. Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37–69, 2003.