Lifetimes of Migration Behavior

Faith Agwang
Computer Science
Department
University of Pretoria
Pretoria, South Africa
faith.agwang@gmail.com

Geoff Nitschke
Computer Science
Department
University of Cape Town
Cape Town, South Africa
gnitschke@cs.uct.ac.za

Willem van Heerden
Computer Science
Department
University of Pretoria
Pretoria, South Africa
wvheerden@cs.up.ac.za

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent agentsMulti Agent Systems

Keywords

Particle Swarm Optimization, Artificial Life

EXTENDED ABSTRACT

Agent Based Modeling (ABM) is a bottom-up approach that has been used to study adaptive group (collective) behavior. ABM is an analogical system that aids ethologists in constructing novel hypotheses, and allows the investigation of emergent phenomena in experiments that could not be conducted in nature [15], [2], [12], [11]. Many studies in ethology have formalized mathematical models of collective migration behavior [1], but few have examined the impact of phenotypic traits (such as lifetime length) on the learning and evolution of collective migration behavior [9], [4].

The first objective of this research is to test the impact of agent lifetime length on the adaptation of collective migration behaviors in a virtual environment. Agent behavior is adapted with a hybrid *Particle Swarm Optimization* (PSO) method that integrates learning and evolution. *Learning* (lifetime learning) refers to a process whereby agents learn new behaviors during their lifetime [13], [3]. *Evolution* (genetic learning) refers to behavioral adaptation over successive lifetimes (generations) of an agent population [5].

The second objective is to demonstrate these hybrid PSO methods are appropriate for modeling the adaptation of collective migration behaviors in an ABM. The motivation is that PSO methods combined with evolution and learning approaches have received little attention as ABMs for potentially addressing (supporting or refuting) hypotheses posited in ethological literature.

The task was for an agent group (flock) to locate a migration point during a *simulated season* in a virtual environment, where a season consisted of X simulation iterations. Thus, varying flock lifetime length (L, where $L \leq X$) varied the number of flock lifetimes (generations) per season. A flock's task performance was measured as the average distance of the flock from the migration point at the end of a season (measured over S seasons).

This study implemented three PSO method variants: *PSO-CA*, *PSO-LT* and *PSO-GT*, extending classical PSO [6] with lifetime and genetic learning, and only lifetime and genetic learning, respectively. PSO [6] with only local best (*lbest*) update and ring neighborhood topology [10] (of a given radius) was used. At each iteration of a flock's lifetime, each particle's fitness was calculated as the inverse of the particle's distance to a migration point. Each particle was initialized with a zero velocity, and could move up to a maximum distance of 0.04 (as as portion of the environment's dimensions) per iteration. All parameter values in this study were derived experimentally, such that minor changes produced negligibly different results for the comparative methods.

At each PSO iteration, each particle's *personal best* fitness (*pbest*) was compared with the current best particle fitness within its neighborhood (*lbest*). If any particle's fitness was greater than *lbest*, then that particle was set as the new *lbest*. The *lbest* only update was used so as to emulate the limited sensory information available to flocking animals in nature.

PSO-LT, PSO-GT and PSO-CA follow this PSO adaptation process but include extensions for lifetime and genetic learning, or both. At each iteration of *PSO-LT*, each particle adopted a current best behavior within its neighborhood. To emulate imperfect copying, each particle copied the behavior of a particle in its neighborhood with a degree of probability proportional to the neighbor's fitness (mimicking fitness proportionate selection [5]). Position and velocity vectors were imperfectly copied, in that a randomly drawn value from a *Cauchy distribution* [8] was added to the particle's position and velocity vectors. If a particle selected itself, then behavioral copying did not occur at that iteration.

Using *PSO-GT* the behaviors of a flock were inherited by the next generation and subject to crossover and mutation operators [5]. PSO-GT thus depended upon multiple flock lifetimes within a season in order for its adaptive process to work. Each particle's genotype directly encoded the position and velocity vector values as an array of floating point values. When a flock's lifetime completed, each particle was selected using fitness proportionate selection [5]. Selected particles were then randomly paired and *1-point crossover* applied. Enough child genotypes were produced to replace the parent population. Mutation with a *Cauchy distribution* was then applied to a randomly selected gene of each child

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.
ACM 978-1-4503-2881-4/14/07.
http://dx.doi.org/10.1145/2598394.2598450.

particle's genotype. The initialization of a child population emulated the birth of a new generation of the flock, where directional information was passed on via the genotypes.

PSO-CA combined the lifetime learning and genetic learning of the *PSO-LT* and *PSO-GT* methods, respectively. At each iteration of *PSO-CA*, after the position and velocity vectors of all particles had been updated, behavioral copying of *PSO-LT* was employed. At the end of a flock's lifetime, the evolutionary operators of *PSO-GT* were applied, and the next generation (lifetime) of the flock initialized.

Experiments applied either *PSO-CA*, *PSO-LT* or *PSO-GT* to collective migration behavior, for flock lifetime lengths in the range [50, 3000]. These lifetime lengths corresponded to [60, 2] generations in one simulated season (3000 simulation iterations). A given experiment was run 30 times and an *average task performance* calculated. The migration point was initialized (in a bounded continuous cubic area) to a random location at the opposite end of the environment to the starting centroid of the flock.

In this study, *short lived* flocks were those with lifetime durations of 50, 100 and 500, and *long lived* flocks were those with lifetime durations of 1000, 1500, and 3000. Results indicated that the *PSO-CA* method, for all short lived flocks, yielded a higher average task performance (with statistical significance indicated by a two-sample t-test [7]) compared to *PSO-LT*, and *PSO-GT* adapted flocks (for any lifetime length). However, all long-lived *PSO-CA*, *PSO-LT*, and *PSO-GT* adapted flocks yielded comparable task performances. The statistically higher task performances of *PSO-CA* adapted (short-lived) flocks is theorized to be a result of the benefit of combining genetic and lifetime learning over a sufficient number of flock generations ([60, 6] lifetimes). For fewer than six lifetimes (lifetime durations: [1000, 3000]), combining genetic and lifetime learning in the context of PSO was not beneficial given the relatively few number of times that evolutionary operators were applied. This resulted in limited opportunity for genetic learning to improve upon good solutions (position and direction vectors) found during a flock's lifetime.

The comparatively lower task performance of *PSO-LT*, and *PSO-GT* (for short lived flocks) is attributed to the lack of global best (*gbest*) information being available to the flock. Since only local neighborhood information (*lbest*) was available to particles many more iterations were required before the flock moved towards the migration point. Also, genetic and lifetime learning were both necessary mechanisms, in that they spread the best solutions throughout the flock thus directing the flock in the correct direction and increasing task performance, given a sufficient number of lifetimes. Hence, long-lived flocks did not outperform short lived flocks (for PSO-LT, PSO-GT, PSO-CA). PSO-CA, combining genetic and lifetime learning, with a sufficient number of lifetimes, demonstrated that both are necessary in order for PSO to effectively adapt collective migration behavior.

Future work entails using more biologically plausible models of the adaptation of collective migration behavior. For example, using agents with developmental genotype to behavioral encodings [14]. Such developmental models will test the environmental and learning conditions under which migration behaviors learned during agent lifetimes are integrated into the genotype. This will further test hypotheses about the role of lifetime duration on genetic versus learned migration behavior [1]. Such future work endeavors to test biological hypotheses pertaining not only to migration behaviors but to broader hypotheses about the transmission of learned behaviors through evolution.

1. REFERENCES

[1] S. Bauer, B. Nolet, J. Giske, J. Chapman, S. Akesson, A. Hedenstrom, and J. Fryxell. Cues and decision rules in animal migration. In *Animal Migration: A Synthesis*, pages 69–87. Oxford University Press, Oxford, UK, 2011.

[2] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems.* Oxford University Press, Oxford, England, 1998.

[3] D. Curran and C. O'Riordan. The effects of cultural learning in populations of neural networks. *Artificial Life*, 13(1):45–67, 2007.

[4] S. Edwards. *The Chaos of Forced Migration: A Means of Modeling Complexity for Humanitarian Ends.* Oxford University Press, Oxford, United Kingdom, 2009.

[5] A. Eiben and J. Smith. *Introduction to Evolutionary Computing.* Springer-Verlag, Berlin, Germany, 2003.

[6] A. Engelbrecht. *Computational Intelligence An Introduction. Second Edition.* John Wiley and Sons, Hoboken, USA, 2007.

[7] B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes.* Cambridge University Press, Cambridge, 1986.

[8] F. Gomez and R. Miikkulainen. Incremental evolution of complex general behavior. *Adaptive Behavior*, 5(1):317–342, 1997.

[9] C. Hemelrijk. The use of artificial-life models for the study of social organization. In *Macaque Societies: A Model for the Study of Social Organization*, pages 295–313. Cambridge University Press, Cambridge, UK, 2004.

[10] J. Kennedy and R. Mendes. Population structure and particle swarm performance. In *Proceedings of the 2002 Congress on Evolutionary Computation*, pages 1671–1676, Honolulu, USA, 2002. IEEE Press.

[11] L. Langenhoven and G. Nitschke. Neuro-evolution versus particle swarm optimization for competitive co-evolution of pursuit-evasion behaviors. In *Proceedings of the IEEE Congress on Evolutionary Computation.* IEEE Press, 2010.

[12] U. Lopez, J. Gautrais, I. Couzin, and G. Theraulaz. From behavioural analyses to models of collective motion in fish schools. *Interface Focus*, 2(2):693–707, 2012.

[13] R. Reynolds. An introduction to cultural algorithms. In *Proceedings of the Third Annual Conference on Evolutionary Programming*, pages 131–139, Singapore, 1994. World Scientifc.

[14] K. Stanley and R. Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.

[15] B. Sumida, A. Houston, J. McNamara, and W. Hamilton. Genetic algorithms and evolution. *Journal of Theoretical Biology*, 147(1):59–84, 1990.