Agents, Pheromones, and Mean-Field Models (Extended Abstract)

H. Van Dyke Parunak
Vector Research Center, a business unit of Jacobs Technology
3520 Green Court, Suite 250
Ann Arbor, MI 48105
+1 734 302 4684
van.parunak@jacobs.com

ABSTRACT

Some agent-based models use digital analogs of insect pheromones for coordination. Such models are intermediate between classical agent-based models and equation-based "mean field" models. Their position in this range can be adjusted by pheromone parameters (notably, the propagation factor).

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence – *multiagent systems.*

General Terms

Algorithms, Design, Experimentation, Theory

Keywords

Modeling, simulation, pheromones, stigmergy, agent interaction

1. INTRODUCTION

Agent-based models (ABMs) and mean-field models (MFMs) have complementary strengths and weaknesses. One approach to ABMs imitates insect pheromones to facilitate coordination. This paper claims that pheromone-based coordination is intermediate between classical ABMs and mean-field models. We confirm this hypothesis with a simple model of population dynamics [4].

2. AGENTS AND MEAN FIELDS

Agent-based models (ABMs) focus on individual entities, while equation-based models (EBMs) focus on variables [3]. EBMs favor global variables, permitting parsimonious closed-form equations. ABMs can use variables accessible to the individual agent, allowing a local viewpoint.

As in statistical physics, a model (EBM or ABM) that replaces individual interactions with system-level averages is a *mean-field model* (MFM). MFMs accept an unrealistic assumption of independence among key variables for improved tractability. In both physics and multi-agent systems, MFMs have limited accuracy [4, 5], but often give more concise insight than discrete models, and researchers often compare both forms of model [1, 2].

The pheromone field in a stigmergic ABM is generated by deposits by individual agents, and is proportional to the probability of encountering an agent at a given location. When an agent makes decisions based on the field, rather than on explicit interaction with other agents, it is reasoning about a weighted average influence of the other agents—weighted because the field is generated by those agents and is concentrated near their locations.

This weighting improves accuracy. Consider five robots in a 20x20 grid. One robot's naïve mean-field estimate of the probability of encountering another robot in any given cell is 4/400 = 0.01. Alternatively, each robot could communicate directly with the others and determine exactly which cells contain other robots. The pheromone approach is intermediate. Each agent contributes to the field locally. The field is an average over agents, localized over limited regions. It is, not a mean-field, but a "lumpy-field."

3. AN EXPERIMENT

A toroidal arena holds two species of agents [4]. Species *I* is immortal, uniformly distributed with average density n_I, and moves randomly with diffusion coefficient D_I. Species *M* is mortal, with initial uniform density n_M. Mortals move randomly with coefficient D_M, die at a constant rate μ, and divide with rate λ when they encounter an immortal. Continuity and symmetry predict that immortals will continue to be homogeneously distributed, $n_I(x) = n_I$. The time evolution of n_M follows

$$\frac{\partial n_M}{\partial t} = D_M \nabla^2 n_M + (\lambda n_I - \mu) n_M \quad (1)$$

For initially uniform spatial distributions of both species, this equation has the time exponential solution,

$$n_M(t) = n_M(0) e^{t(\lambda n_I - \mu)} \quad (2)$$

If $\lambda n_I < \mu$, mortals become extinct.

An ABM without pheromones shows very different behavior. Even for positive values of $\mu - \lambda n_I$ (e.g., 0.3), the mortal population can explode. The difference is due to a mean-field assumption in n_I. As sampled by mortals, immortals are highly nonhomogeneous. Mortals are born next to an immortal. A newly-born mortal sees a local density of immortals far greater than n_I. Some immortals form the core of breeding clusters that generate mortals faster than they can die off.

Whether or not a run with $\lambda n_I < \mu$ explodes depends on stochasticity and location in parameter space. The system parameters λ, μ, D_I, and D_M guide *stochastic* choices by each agent. E.g., a mortal meeting an immortal decides whether to reproduce by uniformly sampling [0, 1] and gives birth only if the result is less than or equal to λ. Different random seeds yield different outcomes. In addition, different *parameters* (population size, birth and death rate, and mortal diffusion rate) affect the persistence of breeding clusters. We observe the effects of these parameters by repeated runs, executing execute a given configuration until mortals either die out or exceed 1000. We repeat each configuration 25 times with different random seeds, and record the percentage of trials in which the mortal population goes to zero.

Cite as: Agents, Pheromones, and Mean-Field Models (Extended Abstract), Parunak, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Yolum, Tumer, Stone and Sonenberg (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1077 - 1078. Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

We also add pheromones, with propagation in space and evaporation in time. In all models, the probability that a mortal gives birth is the product of birth rate λ and the probability *p*(*parent*) that an immortal parent is present. The models differ in how they estimate *p*(*parent*). The MFM estimates $p(parent) = n_I$. In the ABM without pheromones, for each immortal in a cell,

$$p(parent) = \begin{cases} 1 \text{ if immortal is in cell} \\ 0 \text{ otherwise} \end{cases} \quad (3)$$

Pheromones use a "lumpy-field" with a single computation and better accuracy than an MFM. Each immortal deposits one unit of pheromone at its location in each time step. The total deposit at each step equals the immortal population. Each mortal samples the field φ at its location, uses it to estimate *p*(*parent*), and computes the probability of birth. If $\varphi > 1$, the mortal behaves as though it encountered $\lfloor\varphi\rfloor$ immortals, plus one more with probability $\varphi - \lfloor\varphi\rfloor$. This computation is much more efficient than interacting individually with each immortal as in (3). With a single deposit, evaporation = propagation = 0, and stationary immortals, we recover the discrete model. By setting the deposit rate to 1 per immortal and evaporation to 0.5, the total pheromone over the arena is constant, and equal to the immortal population. If immortals do not move and propagation = 0, this configuration also mimics the discrete model. When immortals move, or propagation > 0, the field extends beyond the immortal's cell. This spreading allows invalid births: a mortal may think it is in the presence of an immortal when in fact it is not, the price one pays for a simpler computation.

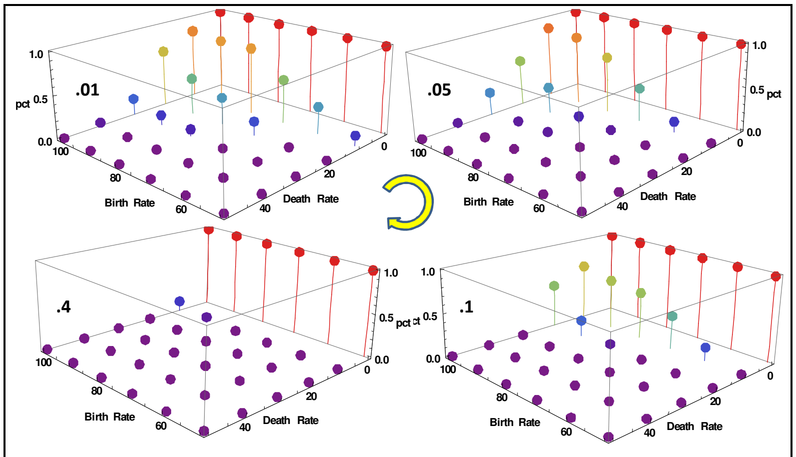

Figure 1: Probability of survival with pheromone propagation of (from top left, clockwise) 0.01, 0.05, 0.1, 0.4

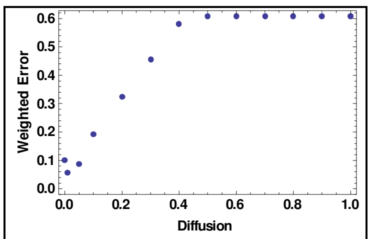

Figure 2: Error vs. diffusion

We model propagation with NetLogo's *diffuse* function, which takes an argument $\rho \in [0,1]$. Each cycle the environment subtracts $\rho * \varphi$ from each cell, and distributes it evenly among the cell's eight neighbors, updating all cells at once.

As ρ increases, a pheromone model should behave less like a discrete model and more like an MFM. However, because the field is stronger near immortals, the error will be less. Figure 1 show this behavior. As ρ increases, probability of survival approaches 0 except when $\mu = 0$, as in the mean-field case.

Each scenario (mean-field and pheromone with various diffusion rates) yields survival rates as a function of birth and death rates that differ from an ABM without pheromones. We weight these differences by the differences from the mean-field case, and normalize by the sum of these weights. On this scale, the MFM scores 1, and the discrete agent system scores 0. Figure 2 shows the variation in this score as a function of propagation. As anticipated, the error grows with propagation rate, and asymptotes before reaching the mean-field level. Unexpectedly, error *increases* when $\rho = 0$, compared with $\rho = 0.01$. $\rho = 0$ corresponds to the discrete model only if the depositing agent is stationary. Our immortals move, leaving a deposit that can mislead mortals. Both propagation and evaporation reduce this obsolete information.

4. CONCLUSION

MFMs avoid the cost of computing individual interactions by replacing them with averages. Conventional ABMs compute each interaction, achieving higher accuracy than a MFM, but the computational burden precludes thorough sampling of the space of possible behaviors.

Pheromones reduce the computational cost of modeling the space of possible interactions, while retaining the interactions of an ABM. The price they pay for this simplification is an approximation. Because the agent framework retains the discrete structure of the problem, the resulting error is often much less than in a complete mean-field treatment, and can be tuned by adjusting the degree of propagation of the pheromones.

Recognizing the mediating position of pheromone models between conventional agents and equation-based MFMs allows modelers to use digital pheromone technology more appropriately.

5. REFERENCES[1]

[1] R. Glinton, P. Scerri, and K. Sycara. Exploiting Scale Invariant Dynamics for Efficient Information Propagation in Large Teams. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2010)*, pages 21-28, IFAAMAS, 2010.

[2] H. V. D. Parunak. A Mathematical Analysis of Collective Cognitive Convergence. In *Proceedings of the Eighth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS09)*, pages 473-480, 2009.

[3] H. V. D. Parunak, R. Savit, and R. L. Riolo. Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users' Guide. In *Proceedings of Multi-agent systems and Agent-based Simulation (MABS'98)*, pages 10-25, Springer, 1998.

[4] N. M. Shnerb, Y. Louzoun, E. Bettelheim, and S. Solomon. The importance of being discrete: Life always wins on the surface. *Proc. Natl. Acad. Sci. USA*, 97(19 (September 12)):10322-10324, 2000.

[5] W. G. Wilson. Resolving Discrepancies between Deterministic Population Models and Individual-Based Simulations. *American Naturalist*, 151(2):116-134, 1998.

[1] The full paper is available at https://activewiki.net/download/attachments/6258699/AAMAS11MeanFieldsFullPaper.pdf