Proceedings of the 2011 Winter Simulation Conference
S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, eds.

INTERACTION METRIC OF EMERGENT BEHAVIORS IN AGENT-BASED SIMULATION

Wai Kin Victor Chan

Department of Industrial and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180, USA

ABSTRACT

Agent-based simulation (ABS) has been a popular tool in various science and engineering domains. Simulating emergent behavior is one main usage of ABS. This paper investigates the use of interaction statistics as a metric for detecting emergent behaviors from ABS. An emergent behavior arises if this interaction metric deviates from normality.

1 INTRODUCTION

Agent-based simulation (ABS) is a fast growing area in operations research and industrial engineering. It is an approach capable of tackling large-scale problems in various sciences and engineering disciplines. It is a tool that has been facilitating collaborations between industrial engineers and researchers in other areas. A survey on practices of ABS is given in (Heath, Hill, and Ciarallo 2009).

One of the main usages of ABS is the simulation of emergent behaviors of underlying complex systems. While there exist many other simulation approaches that can be used to simulate emergent behaviors at various levels of details, ABS stands out from existing approaches because of its flexibility in simulating interacting autonomous agents; and interactions give rise to emergence behaviors. This paper focuses on this usage of ABS. We first review the concept of emergence and existing work on emergence.

The study of emergent behaviors has been a journey of more than 130 years long, probably set sailed at the work by (Lewes 1875). It is certainty an intriguing topic for otherwise it would not have lasted for that long and many questions including the most fundamental one—what is emergence?—would have been solved. Readers interested in the historical and philosophical roots of emergence can consult with (Goldstein 1999). Other good resources of the history of emergence are (Corning 2002; Johnson 2006).

Emergent behaviors can be found almost everywhere. One quick example is the human mind, which emerges from a large number of neurons collaborating together in a not fully understood way (Bedau 1997). (As a side note, while many people are fascinated about the human mind, no one knows whether the neurons are really collaborating in the optimal way or simply following some universal law, for example, would there be another way of neuron collaboration that will produce emergence superior in functionalities to our current mind?) The formation of human society is another example of emergence. Recent development of social networking has enabled people to interact online to carry out businesses, information exchange, or simply networking—all these can be considered as new emergent phenomena that could not otherwise be done without the human interactions facilitated by the Internet (Haglich, Rouff, and Pullum 2010). Last but not least, the flocking of birds, schooling of fishes, or herding of animals, are all examples of emergent behaviors.

Emergent behavior, by its name, is a behavior or pattern emergent (i.e., supervenient) from its constituents (or parts). Whether this "emergence" is traceable to its constituents is a question that has flourished a decade-long debate and resulted in several definitions of emergent behaviors. Indeed, the identification of emergence is inherently subjective (see discussion in Crutchfield 1994). (Chan, Son, and Macal 2010) call a behavior emergent if it is interesting to at least one observer, for otherwise, the identification of the behavior provides no societal value. This definition involves an observer. Definitions of emergence that are independent of the observer can be found in (Crutchfield 1994).

Traditional definitions of emergence emphasize on the unpredictability of emergent behaviors. Proponents believe that emergent behaviors must be unexpected, for otherwise they are not emergent, and cannot be reduced to its constituents (Lewes 1875). In other words, the behavior is not owned by any one of its constituents nor even by all constituents. The behavior only appears when all constituents interact in a right manner. The saying "the whole is larger than the sum of its parts" summarizes the essence of this line of definitions. Behaviors of this sort are described as emergence rather than resultant.

Others criticize that the traditional definitions are metaphysical and inherently problematic or scientifically irrelevant due to their support of "irreducible downward causation" (Bedau 1997). Bedau introduces a notion of weak emergence to distinguish emergent behaviors observed in simulation from metaphysical emergences. A weakly emergent behavior is a behavior obtainable only by simulation. Because they are derivable by simulation, weak emergences are predictable, at least in theory, and therefore, scientifically useful. Weak emergence not only materializes the concept of emergence but also scientifically casts emergence as a consequence of computation—i.e., simulation. Pros and cons of weak emergence are listed in (Baker 2010). One main critic is that weak emergence shifts the difficulty of defining emergence to the definition of "simulation." Baker distinguishes simulation from analytical derivation (i.e., mathematical analysis) and shows an example in which analytical derivation is needed in determining emergence.

(Corning 2002) explains emergence in terms of the synergy of parts. Corning argues that synergies are abound but not all are emergence. For a behavior to be qualified as emergence, underlying parts must have certain effects to each other (like reshaping each other) when they interact and participate in the whole. According to (Corning 2002), water and human body are examples of emergence as molecules are collaborating and shaping each other's properties to make the whole. On the other hand, as argued in (Corning 2002), a sand pile or a river do not qualify as emergent phenomena.

Classifications and taxonomies of emergent behaviors have also been proposed, see, for example, (Bar-Yam 2004; Fromm 2004; Gore and Reynolds Jr 2007).

One main debate about emergence is whether or not an emergent behavior is derivable from its parts. Some support the underivable answer because the behavior is not a property of the parts. Others respond that it is underivable because the behavior is a property of the whole when the parts act together. Still, some reject this claim, arguing that the parts are acting in such a complex way that downward causality is impossible to trace. Reconcilers then propose categorizing emergent behaviors into derivable (e.g., via computation) and underivable (e.g., metaphysical) so that the society can "derive" values from the derivables, which are scientifically researchable.

Is it really underivable in principle? Or is it underivable because of the limitation of present techniques? Or is it underivable due to some paradox like the observation paradox or quantum uncertainty, that is, measurements or observations inevitably influence the parts being observed and thus actuate observation is impossible? Rather than trying to tackle this long standing issue, this paper will focus on the derivable emergent behaviors, i.e., those observed from a simulation and every details can be pinpointed.

For decades, people have been trying to measure or quantify emergence. A summary of tools for detecting emergent behavior is given in (Boschetti et al. 2005). Various mechanisms have been proposed as metrics for detecting or evaluating emergence. Information is one of such measures. Using information entropy, (Wuensche 1999) demonstrates that complex patterns of one-dimensional cellular automata exhibit high-variance of input-entropy over time. (Crutchfield 1994) introduces a hierarchical framework based on computational mechanics to study emergence and proposes complexity metrics that take into account deterministic complexity and stochastic complexity to detect and quantify emergence. (Haglich, Rouff, and Pullum 2010) uses semi-Boolean algebra to detect emergent behaviors in social networks.

In work related to data mining, various methods have been introduced to find patterns or changing points from a set of observed data. If the data is the observations (e.g., samples over time) of a multi-agent system, then existing methods for detecting patterns, changing points (e.g., phrase transition), or even outliers (e.g., for validation) can be applied (some may need modifications) to detect emergent behaviors. For instance, (Shalizi 2001) gives an algorithmic approach for pattern discovery in time series and cellular automata. (Grossman et al. 2009) define emergent behavior as a changing point in a time-series and suggest that changing point detection algorithms can be used to detect emergent behaviors.

In the fields of social science, computer science, and robotics, there are also a large amount of work done in studying emergent behaviors. For example, (Minati 2002) uses ergodicity to detect emergence. Ergodicity in (Minati 2002) is different from that of queueing systems. In (Minati 2002), ergodicity means that the average behavior of a set of agents at a particular time epoch is similar to the average behavior of an individual agent of this set over a long period of time. (For readers familiar with queueing systems, one analogy would be the PASTA property of a Markovian queue—Poisson Arrivals See Time Averages (Wolff 1982).) Ergodicity in a multi-agent system changes when emergent behavior arises; the ergodicity of the whole system may diminish with the raises of ergodicity in different sets of agents who are forming local emergent behaviors (such as patterns in Game of Life).

(Gore and Reynolds 2008) study emergent behaviors based on causal relationships among events in a simulation. (Hovda 2008) quantifies emergence based on "the amount of simulation" needed to obtain the behavior. (Schaefer et al. 2002) introduce a meta-architecture (e.g., based on UML) that maps inputs and outputs of a simulation so as to identify emergent behaviors from the outputs.

While there have been quite an amount of work done on emergence and various metrics have been proposed, the basic ingredient of emergence—interactions—has not been explicitly estimated and used as a metric to detect emergence behaviors. We investigate the use of the metric—interaction counts—in detecting emergent behaviors. We apply and test this metric on three ABS models: the Game of Life model, the Boids model, and the Brownian motion model.

The reminder of the paper is organized as follows. Section 2 defines the interaction metric in each of the three ABS models. Section 3 presents experimental results. Section 4 makes a conclusion.

2 INTERACTION METRIC

While there are still many unknowns in emergence, one thing for sure is that no emergence can arise without interactions from its parts. Interaction is a necessary key. Therefore, in this paper we examine the use of interactions in identifying emergent behaviors.

In dictionary, interaction is defined as "a mutual or reciprocal action or influence." This definition concerns more about the outcome from an interaction rather than the process of the interaction. Similarly, Wiki describes an interaction as "a kind of action that occurs as two or more objects have an effect upon one another" (as of June 2011). To facilitate our study, we define interactions as follows.

This paper is within the scope of ABS. In this paper (under the ABS sense), an interaction occurs when an agent initiates or receives a contact with another agent. The contact can be of any form, including information query or "physical" contact in a computer-model based environment (e.g., billiard ball collisions).

In a high resolution model, an interaction always triggers state changes. For instance, although an agent after interacting with another one may decide to do nothing, this interaction does trigger the agent to go through a decision process, which can be considered as a state change even thought this change disappears in the end. For other models in which the decision process is not considered as a state change, an interaction may or may not trigger a state change. As a result, depending on the objective of a study, interactions may or may not be measured by state changes. We define two types of interactions in the following:

- Regular interaction: A regular interaction occurs whenever an agent initiates or receives a contact with another agent regardless of whether this action will finally induce any outcome.
- Effective interaction: An effective interaction is a regular interaction with a final outcome.

Interactions are mutual. But when counting the number of interactions, it helps if we distinguish directional interactions from undirectional interactions. Directional interactions are interactions initiated by one of the two interacting agents. Undirectional interactions are interactions happened spontaneously between two agents or caused by a third party or other forces not belonging to the two interacting agents. For example, in the Boids model, an agent must initiate inquires to figure out its nearest agent. In this case, the interactions are directional: from the agent making the inquiries to the surrounding agents being inquired. In the Brownian motion model, when two agents collide, an undirectional interaction occurs. This is undirectional because it is a result of two agents getting too close to each other. Undirectional interactions are usually double counted (i.e., both agents' interaction counters are incremented) since it is hard to attribute the interaction to one of the two interacting agents. On the other hand, it is the modeler's decision of whether to double count or single count the directional interactions. In the two models below (Game of Life and Boids), we will single count the directional interactions. That is, only the interaction counter of the initiating agent is incremented.

We also note that other definitions of interactions are possible. But to facilitate our study, we shall stick to these two definitions above.

The interaction measure should be easy to implemented and computed. We use the ABS execution algorithm given in (Chan, Son, and Macal 2010) (but remove the continuous variables) to illustrate where in a typical ABS model to add code to keep track of the interaction statistics. This algorithm is repeated in Figure 1. In most cases, this piece of bookkeeping code can be added right after an agent has finished interacting with other agents. The bold face words shown in Figure 1 give an example of where to add such piece of code.

Because an interaction can be different forms in different models, in the following, we shall first define how we measure the interactions in each of the three models studied in this paper.

<u>Agent-Based Simulation Execution Algorithm with Continuous State Variables:</u>

- Initialize
- Do until stop condition is satisfied
- For each agent
    - Perform actions: change state, send or fetch messages to (or from) other agents or the environment, etc.
    - **If this agent has interact with other agent(s), increment the interaction counter by the number of agents with which it has interacted.**
- **Increment the global counter for interactions of all agents.**
- Advance clock

Figure 1: An ABS Execution Algorithm without Continuous State Variables

2.1 Game of Life

The first model is Conway's Game of Life model (Berlekamp, Conway, and Guy 2003). Imagine that there is a two-dimensional grid of size $n \times n$, where n is the number of rows (or columns) in the grid. A cell lives in each entry of the grid. The cells cannot move. Each cell has only two states: alive or dead. Each cell changes its state based on three simple rules (see below) that describe the interactions between a cell and its eight neighbors (up, down, left, right, and four diagonal cells). Cells along the edges have five neighbors and cells at the four corners have only three neighbors. Figure 4(a) shows a screen shot of this model at time 200. A colored dot (or dark in monochrome) represents a live cell. Dead cells are uncolored and therefore, appear as white (or empty) in the figure.

Three agent interaction rules in the Game of Life model are:

1. A live cell with exactly 2 or 3 live neighbors will remain alive in the next time step.
2. A dead cell with exactly 3 live neighbors will come to life in the next time step.
3. Otherwise, the cell will die either of loneliness or overcrowding in the next time step.

The flexibility of this model in arriving at different emergent patterns has been demonstrated by many people, see, e.g., (Rendell 2002) for an extensive discussion of various emergent patterns. Various extensions of the Game of Life models have also been studied. (Chan 2010) generalizes the Game of Life model by changing the number of live neighbors in Rules #1 and #2 to X, Y, and Z, which are defined below. This variation has a total of 405 possible combinations of rules, (X, Y, Z). Some of them are trivial but some produce clear emergent patterns. One can also use a series of rules to obtain different combinations of patterns. See (Chan 2010) for examples of these patterns. We call this model the variation of Game of Life model and the three rules are:

1. A live cell with at least X and at most Y live neighbors will remain alive in the next time step.
2. A dead cell with exactly Z live neighbors will come to life in the next time step.
3. Otherwise, the cell will die either of loneliness or crowdedness in the next time step.

where $0 \leq X, Y, Z \leq 8$ and $X \leq Y$.

In each iteration, because every cell must check the status of its neighbors, this amounts to eight inquiries (or equivalently, eight regular interactions) carried out by an agent. Agents along the edges [at the corners] initiate five [three] regular interactions. All interactions here are directional and we only count the interactions for the initiating agents. Moreover, these numbers of interactions are deterministic. Precisely, the total number of interactions at each iteration is equal to the sum of neighbors of all cells, i.e., $8 \times (n-2)^2 + 5 \times 4 \times (n-2) + 3 \times 4 = 8n^2 - 4n + 4$. Because of this deterministic number, we switch our attention to the effective interactions, i.e., the number of state changes caused by these deterministic interactions.

Specifically, we count the number of cells that change their state at each iteration. Suppose the simulation is run for *T* iterations. We define the following notation to facilitate the discussion. Let $s_{it}, i = 1, \ldots, n^2, t = 1, \ldots, T$, be the state of cell *i* at time *t* with s_i = 1 if alive and s_i = 0 if dead. Let δ_{it} be the indicator function:

$$\delta_{it} = \begin{cases} 1, & s_{it} \neq s_{i,t-1} \\ 0, & s_{it} = s_{i,t-1} \end{cases}$$

Let $I_t = \sum_{\forall i} \delta_{it}$ be the total state changes at time *t*. We also investigate the cumulative state changes. Let X_{it} be the number of state changes of cell *i* from time 1 until time *t*, i.e., $X_{it} = \sum_{l=1,\ldots,t} \delta_{il}$. The cumulative state changes are normalized by the maximum cumulative state changes. Let $X_t^* = \max\{X_{it}; i = 1, \ldots, n^2\}$ be the maximum cumulative state changes among all cells. The normalized cumulative state changes for cell *i* is $Y_{it} = X_{it}/X_t^*$. The total cumulative state changes up to time *t*, denoted by Z_t, is the sum of all individual cumulative state changes, i.e., $Z_t = \sum_{\forall i} X_{it}$. I_t is in fact the finite derivative of Z_t, as shown below:

$$Z_t - Z_{t-1} = \sum_{\forall i} X_{it} - \sum_{\forall i} X_{i,t-1} = \sum_{\forall i} \sum_{l1,\ldots,t} X_{il} - \sum_{\forall i} \sum_{l=1,\ldots,t-1} X_{il} = \sum_{\forall i} X_{it} = I_t$$

We will examine the following statistics in the next section:

1. Time series: $I_t, t = 1, \ldots, T$
2. Distribution of $Y_{it}, t = 1, \ldots, T$
3. Time series: $Z_t, t = 1, \ldots, T$

2.2 Boids

The Boids model, introduced by (Reynolds 1987), simulates the flocking behavior of birds (or schooling behavior of fishes) emergent from leaderless interactions among self-propelled agents (i.e., fishes or birds). The emergent flocking behavior is also similar to the grouping behavior of other creatures, such as schooling behavior in chub mackerel, *Scomber japonicus*, and herding behavior in animals and crowds.

In the Boids model, each agent, or boid, follows three simple rules:

4. Separation: steer to avoid crowding local flockmates
5. Alignment: steer towards the average heading of local flockmates
6. Cohesion: steer towards the average position of local flockmates

The pseudocode of the model is given in the following.

<u>Boids Model:</u>

- Define boid agents, parameters, and variables
- Initialize: Create boids.
- Do until stop
  - For each boid:
    - Find Flockmates within vision distance.
    - If Flockmates are not empty
      - Determine the Closest boid in Flockmates
    - If the Closest boid is too close, then
      - turn away from the Closet boid (Separate).
    - Else
      - turn towards the average heading of Flockmates (Align), and
      - turn towards the average position of Flockmates (Cohere).
- Advance clock by one tick

Figure 2: Boids Model Pseudocode

In the Boids model, an interaction occurs when an agent is within a distance threshold from its nearest agent. This distance threshold is called the vision of an agent. Let v be this distance threshold, which for simplicity is assumed to be the same for all agents. Nearest agents are not necessary mutual; that is, if Agent A is the nearest agent of Agent B does not mean Agent B is also the nearest agent of Agent A. There could be another Agent C that is nearer to Agent A than Agent B does, see Figure 3 below. The interactions here are directional and only the counters of the initiating agents are incremented when interactions occur.

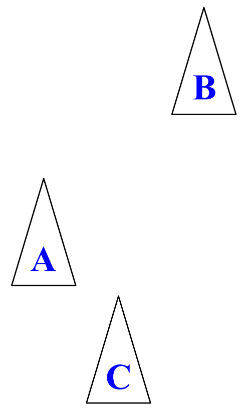

Figure 3: Nearest Neighbors in Boids Model

The outcome of an interaction is either the action of separation or the actions of alignment and cohesion. Once an agent flocks with other agents, interactions take place repeatedly unless the agent joins another flocking group.

Let $d_{ijt}, i \neq j = 1, \ldots, n^2, t = 1, \ldots, T$, be the distance from agent *i* to agent *j* at time *t*. Let $d_{it}^{\min}, i = 1, \ldots, n^2, t = 1, \ldots, T$, be the minimum distance between agent *i* to any other agent $j \neq i$ at time *t*. Let δ_{it} be the indicator function:

$$\delta_{it} = \begin{cases} 1, & d_{it}^{\min} \leq v \\ 0, & d_{it}^{\min} > v \end{cases}$$

The state variables, $I_t, X_{it}, X_t^*, Y_{it}$, and Z_t are defined in the same manner.

2.3 Brownian Motion

Brownian motion describes the random movement of a particle suspended in liquid, where the particle is subject to a large number of small molecular shocks. This physical phenomenon has been well studied (Brown 1828; Einstein 1906; Nelson 1967). These studies of Brownian motion have made a profound impact on many science domains. For instance, Brownian motion proved the existence of atoms, enabled us to measure precisely their sizes, redefined the theory of thermodynamics, created a new branch of physics on fluctuation phenomena, and laid the foundations for statistical thermodynamics and the theory of stochastic processes.

Here, we examine the overall interactions of molecules from a simulation point of view. All molecules and the particle are modeled as billiard balls. Initially, they are randomly distributed on a two-dimensional box. Their initial travel directions are also randomly generated. Their size and initial speed (which determine their energy) are parameters set by users.

As the simulation runs, the agents move inside a two-dimensional environment (see Figure 9(a)). When two agents collide, they exchange energy in accordance to the law of physics (i.e., elastic collisions). Therefore, the rules governing the interactions between the agents are the simple physical laws of elastic collision, in which both momentum and kinetic energy are conserved.

The previous two models execute based on iterations. Each iteration is equivalent to one unit of time in the model. Unlike these two models, the BM model executes based on the occurrences of collisions. To do that, it maintains a list of future collisions (i.e., the time epochs of the collisions and the names of the colliding agents) and advances its clock one collision by one collision. The list is updated if a collision causes new future collisions or cancels some previous scheduled collisions.

An interaction occurs when two agents (e.g., two molecules or a molecule and the particle) collide with each other. Because the interaction is undirectional, both agents will increase their interaction counters. Let δ_{it} be the indicator function:

$$\delta_{it} = \begin{cases} 1, & \text{if collides at time } t = \{t_1, t_2, \ldots, t_k, \ldots\} \\ 0, & \text{o.w.} \end{cases}$$

where t_k is the time epoch at which the k^{th} collision occurs. Because only one collision can occur at each time, $I_t = \sum_{\forall i} \delta_{it}$ is always 2 (1 for each of the two interacting agents). Other state variables, X_{it}, X_t^*, Y_{it}, and Z_t are defined similarly as in the previous two subsections.

3 EXPERIMENT

Various emergent behaviors (or patterns) have been found in the Game of Life, such as blocks, blinkers, and gliders. These patterns can be observed in Figure 4(a), which shows a screen shot of the model with rule (*X*, *Y*, *Z*) = (2, 3, 3) at time 200 (the original Game of Life model). Figure 4(b)-(d) show the time series of I_t, the histogram of Y_{it}, and the time series of Z_t. The fast decrease in I_t is due to a transient period from the initial random distribution of live and dead cells to the state of several sustainable emergent patterns. The patterns are maintained by interactions of small local cells and these interactions contribute to some slight oscillations in I_t, which look periodic. The histogram of the normalized cumulative interaction (Y_{it}) becomes more and more skew with the right tail representing the increasing interactions of the local cells that sustaining the emergent patterns. The histogram is similar to a power-law distribution in the long run. It is not difficult to foretell the continuous increase in the total cumulative interactions (Z_t). What deserves a note here is that Z_t is concave in *t* when emergent patterns appear. We will see that Z_t will tend to be a straight line if no emergent pattern appears in the later experiments.

We next examine the variation Game of Life model with the rule (*X*, *Y*, *Z*) = (4, 7, 5). This rule leads to either a fix group of live cells (i.e., a fixed pattern) or zero live cell (i.e., an empty world) depending on the initial random distribution of live cells. (Simulating this rule for 30 times resulted in 3 runs with zero live cell.) One example of fixed pattern is shown in Figure 5(a). Also, from Figure 5(b)-(d), we see that I_t decreases to zero because the pattern is fixed and no cell changes its state, the histogram of Y_{it} shows the frequency of state changes of all cells from the initial state to the fixed pattern, and Z_t quickly increases to a fixed value representing the total number of state changes needed to obtain the fixed pattern.

Both Figure 4 and Figure 5 show emergent patterns (both evolving in time or fixed). The statistics, therefore, do not appear as normal, e.g., Figure 4(c) and Figure 5(c) do not following a normal distribution.

We now study the rule (*X*, *Y*, *Z*) = (6, 7, 1), which does not produce a clear pattern like those in Figure 4 and Figure 5; it produces random distributions of live cells at each iteration. Figure 6(a) gives a screen shot of this model. The distribution of live cells changes chaotically. I_t oscillates at a high value (Figure 6(b)); Y_{it} appears to follow a normal distribution (Figure 6(c)); and Z_t increases linearly in time (Figure 6(d)). The normality of Y_{it} suggests the absence of recognizable emergent patterns.

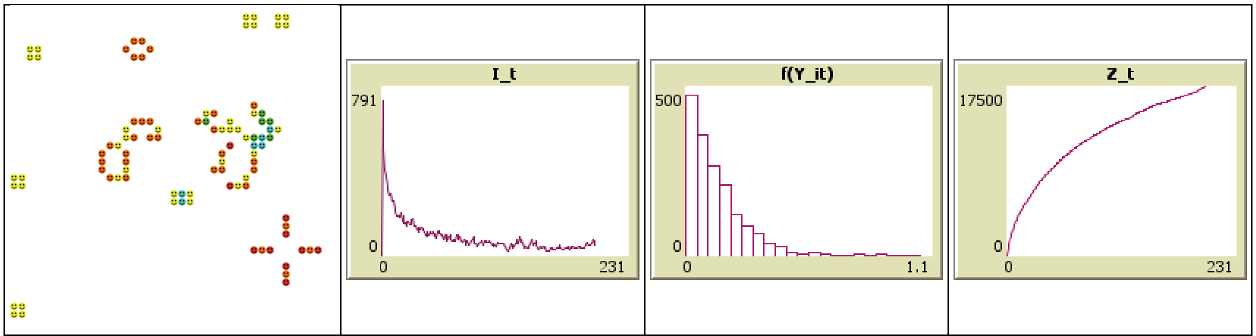

(a) Screen Shot | (b) I_t | (c) Y_{it} | (d) Z_t

Figure 4: Game of Life: (X, Y, Z) = (2, 3, 3)

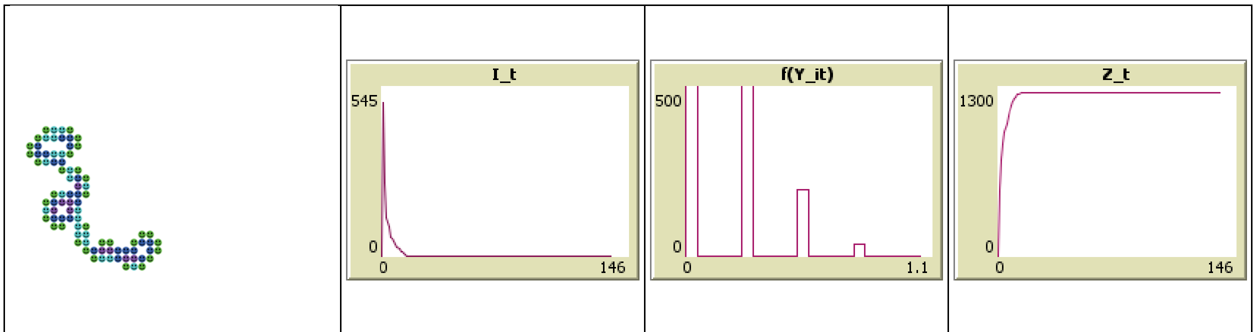

(a) Screen Shot | (b) I_t | (c) Y_{it} | (d) Z_t

Figure 5: Game of Life: (X, Y, Z) = (4, 7, 5)

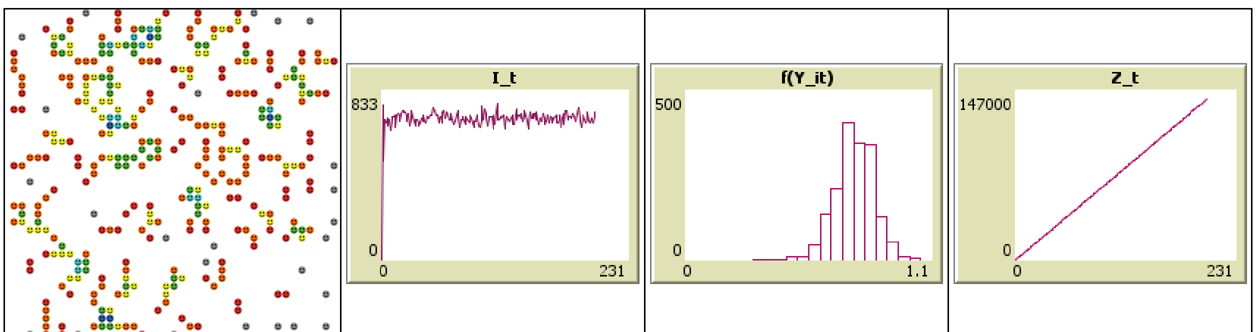

(a) Screen Shot | (b) I_t | (c) Y_{it} | (d) Z_t

Figure 6: Game of Life: (X, Y, Z) = (6, 7, 1)

Emergent patterns in the Boids model can be easily identify visually—emergent behavior is observed when flocking behavior appears. Figure 7(a) shows a screen shot of the original Boids model at time (iteration) 2012. Several flocking groups are observed. The flocking of agents maintains the level of their interactions, i.e., when an agent flock with other agents, it continues to interact with the closest agent unless it flies away from this group. Therefore, I_t converges to n—the number of agents, which is 150 in the example shown in Figure 7. Z_t as the integral of I_t first experiences a short shape (convex) increase and then increases linearly in time due to the near fixed value of interactions when I_t converges to n. The histogram of Y_{it} was initially quite uniform and eventually converged to a spike on the right end, showing a clear deviation from normality. This right-tailed spike is because the interactions of all agents equalize when all of them flock for a sufficient amount of time.

There are many possible ways of destroying the flocking behavior, either by altering the parameters or by changing the rules. Here, we simply remove the alignment rule and only keep the separation and cohesion rules. Figure 8(a) shows that after 2000 iterations, no flocking behavior is observed. Just like the model in Figure 6, I_t resembles a random walk (because of the irregular behavior) with Z_t increasing linearly and Y_{it} converging to a distribution similar to a normal distribution.

Finally, we examine a model that is inherently normal and see if the interaction statistics also exhibit normality. The Brownian motion model is shown in Figure 9(a). As seen in Figure 9(c), the distribution of Y_{it} resembles a normal distribution, suggesting the absence of emergent behavior. The plots of I_t and Z_t also match our expectation, i.e., taking value of 2 and increasing linearly, respectively.

All the models were created in NetLogo (Wilensky 1999).

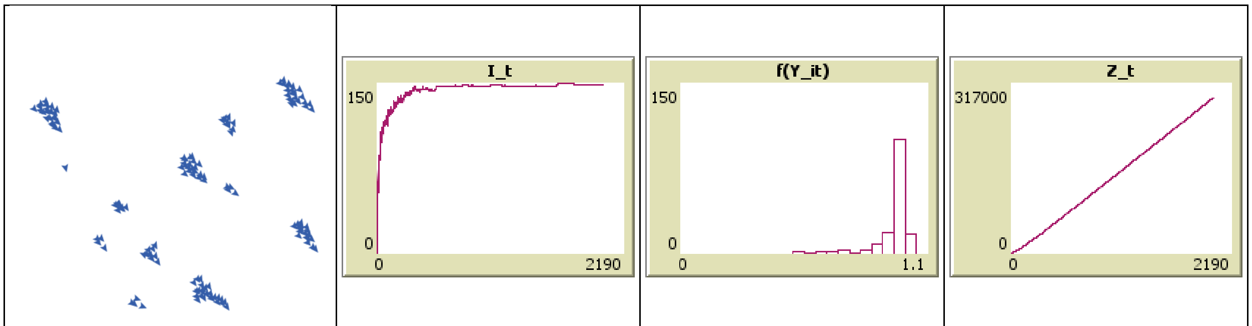

(a) Screen Shot | (b) I_t | (c) Y_{it} | (d) Z_t

Figure 7: Boids: (Separation, Alignment, Cohesion)

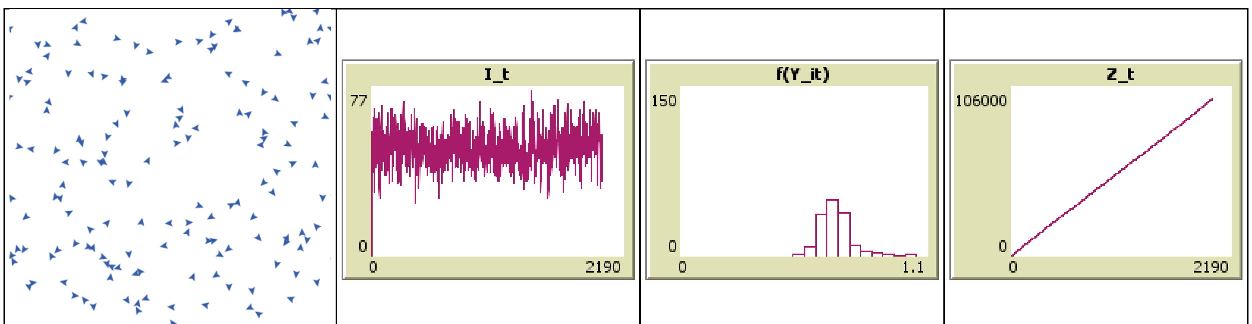

(a) Screen Shot | (b) I_t | (c) Y_{it} | (d) Z_t

Figure 8: Boids: (Separation and Cohesion, No Alignment)

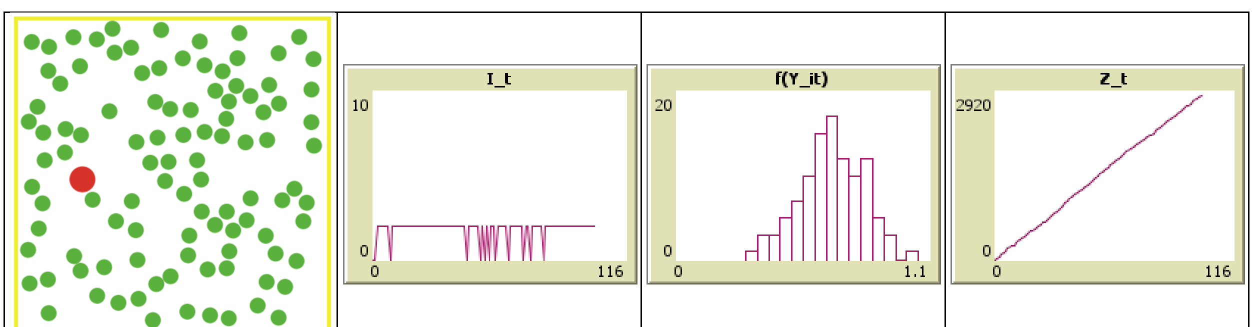

(a) Screen Shot | (b) I_t | (c) Y_{it} | (d) Z_t

Figure 9: Brownian motion Model

4 CONCLUSION

This paper investigates the use of interactions as a metric for detecting emergent behaviors in three ABS models: Game of Life, Boids, and Brownian motion. In both the Game of Life and Boids models, when emergent behaviors arise, the interaction metric deviates from normality. On the other hand, when emergent behavior is absent, the interaction metric behaves in accordance to a distribution similar to normal. This result is confirmed in the Brownian motion model. The Brownian motion model only simulates random movements of particles and inherently does not exhibit emergent behaviors without external perturbations. It is found that the interaction metric distributes quite normally. This suggests that interactions could be a metric for detecting emergent behaviors, at least for these three examples.

We make no claim that the interaction metric is valid for all ABS models. More studies are needed to evaluate this metric when used for different ABS models.

REFERENCES

Baker, A. 2010. "Simulation-based definitions of emergence." *Journal of Artificial Societies and Social Simulation*, 13(1):9.

Bar-Yam, Y. 2004. "Multiscale variety in complex systems: Research articles." *Complex*., 9(4):37-45.

Bedau, M. 1997. "Weak emergence." In *Philosophical perspectives: Mind, causation, and world*. ed. J. Tomberlin. 11**:**375-399: Blackwell Publishers.

Berlekamp, E., J. Conway, and R. Guy. 2003. *Winning ways for your mathematical plays.* 2nd. Natick, Massachusetts: AK Peters, Ltd.

Boschetti, F., M. Prokopenko, I. Macreadie, and A.-M. Grisogono. 2005. "Defining and detecting emergence in complex networks." In *Knowledge-based intelligent information and engineering systems***:**573-580.

Brown, R. 1828. "A brief account of microscopical observations made in the months of June, July and August, 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies." *Phil. Mag*., 4:161-173.

Chan, W. K. V. 2010. "Foundations of simulation modeling." In *Encyclopedia of operations research and management science*. ed. J. J. Cochran. New York: Wiley.

Chan, W. K. V., Y.-J. Son, and C. M. Macal. 2010. "Agent-based simulation tutorial - simulation of emergent behavior and differences between agent-based simulation and discrete-event simulation." In *Proceedings of the 2010 Winter Simulation Conference.* eds. B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan. 135-150. Piscataway, New Jersy: Institute of Electrical and Electronics Engineers.

Corning, P. A. 2002. "The re-emergence of "Emergence": A venerable concept in search of a theory." *Complexity*, 7(6):18-30.

Crutchfield, J. P. 1994. "The calculi of emergence: Computation, dynamics and induction." *Physica D: Nonlinear Phenomena*, 75(1-3):11-54.

Einstein, A. 1906. "A new determination of molecular dimensions." *Annalen der Physik*, 4(19):289-306.

Fromm, J. 2004. *The emergence of complexity.* Kassel: Kassel University Press.

Goldstein, J. 1999. "Emergence as a construct: History and issues." *Emergence: Complexity and Organization*, 1(1):49-72.

Gore, R., and P. Reynolds Jr. 2007. "An exploration-based taxonomy for emergent behavior analysis in simulations." In *Proceedings of the 2007 Winter Simulation Conference.* eds. S. G. Henderson, B. Biller, M.-H. Hsieh, J. Shortle, J. D. Tew, and R. R. Barton. 1232-1240. Piscataway, New Jersy: Institute of Electrical and Electronics Engineers.

Gore, R., and P. Reynolds. 2008. "Applying causal inference to understand emergent behavior." In *Proceedings of the 2008 Winter Simulation Conference.* eds. S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, and J. W. Fowler. 113-122. Piscataway, New Jersy: Institute of Electrical and Electronics Engineers.

Grossman, R., M. Sabala, Y. Gu, A. Anand, M. Handley, R. Sulo, and L. Wilkinson. 2009. "Discovering emergent behavior from network packet data: Lessons from the angle project."

Haglich, P., C. Rouff, and L. Pullum. 2010. "Detecting emergence in social networks." In *Proceedings - SocialCom 2010: 2nd IEEE International Conference on Social Computing, PASSAT 2010: 2nd IEEE International Conference on Privacy, Security, Risk and Trust.* 693-696. IEEE Computer Society.

Heath, B., R. Hill, and F. Ciarallo. 2009. "A survey of agent-based modeling practices (January 1998 to July 2008)." *Journal of Artificial Societies and Social Simulation*, 12(4):9.

Hovda, P. 2008. "Quantifying weak emergence." *Minds Mach.*, 18(4):461-473.

Johnson, C. 2006. "What are emergent properties and how do they affect the engineering of complex systems?" *Reliability Engineering and System Safety*, 91 12:1475-1481.

Lewes, G. H. 1875. *Problems of life and mind.* London: Londan Kegan Paul, Trench, Trübner.

Minati, G. 2002. "Emergence and ergodicity: A line of research." In *Emergence in complex cognitive, social and biological systems*. eds. G. Minati, and E. Pessa. New York, NY: Kluwer Academic Publishers.

Nelson, E. 1967. *Dynamical theories of brownian motion.* Princeton, NJ: Princeton University Press.

Rendell, P. 2002. "Turing universality of the game of life." In *Collision-based Computing*:513-539: Springer-Verlag.

Reynolds, C. 1987. "Flocks, herds, and schools: A distributed behavioral model." *SIGGRAPH '87*, Computer Graphics (ACM) 21(4):25-34.

Schaefer, L., P. Wolfe, J. Fowler, and T. Lindquist. 2002. "Simulation meta-architecture for analyzing the emergent behavior of agents." (Draft Manuscript: http://www.csu.edu.au/ci/draft/schaef01).

Shalizi, C. 2001. "Causal architecture, complexity and self-organization in time series and cellular automata." University of Michigan.

Wilensky, U. 1999. Netlogo. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University. Available via <http://ccl.northwestern.edu/netlogo/> [accessed October 10, 2010].

Wolff, R. 1982. "Poisson arrivals see time averages." *Operations Research*, 30(2):223-231.

Wuensche, A. 1999. Classifying cellular automata automatically: Finding gliders, filtering, and relating space-time patterns, attractor basins, and the z parameter. Complexity 4(3):47-66.

AUTHOR BIOGRAPHIES

WAI KIN (VICTOR) CHAN is an Associate Professor of the Department of Industrial and Systems Engineering at the Rensselaer Polytechnic Institute, Troy, NY. He holds a Ph.D. in industrial engineering and operations research from University of California, Berkeley. His research interests include discrete-event simulation, agent-based simulation, and their applications in energy markets, social networks, service systems, transportation networks, and manufacturing. His e-mail address is <chanw@rpi.edu>.