Benjamin D. Nye*
Barry G. Silverman
Ackoff Collaboratory for Advancement of the Systems Approach (ACASA),
University of Pennsylvania
Philadelphia, PA 19104

Social Learning and Adoption of New Behavior in a Virtual Agent Society

Abstract

Social learning and adoption of new behavior govern the rise of a variety of behaviors: from actions as mundane as dance steps to those as dangerous as new ways to make IED detonators. However, agents in immersive virtual environments lack the ability to realistically simulate the spread of new behavior. To address this gap, a cognitive model was designed that represents the well-known socio-cognitive factors of attention, social influence, and motivation that influence learning and the adoption of a new behavior. To explore the effectiveness of this model, simulations modeled the spread of two competing memes in Hamariyah, an archetypal Iraqi village developed for cross-cultural training. Diffusion and clustering analyses were used to examine adoption patterns in these simulations. Agents produced well-defined clusters of early versus late adoption based on their social influences, personality, and contextual factors, such as employment status. These findings indicate that the spread of behavior can be simulated plausibly in a virtual agent society and has the potential to increase the realism of immersive virtual environments.

1 Introduction

Virtual environments are approaching a paradigm shift from virtual agents to virtual agent societies. This is a transition toward rich modeling of the interactions between virtual agents, rather than just agent–user interaction and agent–environment interaction. This shift has already started in industry and research applications. The next step is to utilize full-fledged virtual agent societies in immersive environments, as used by training, teaching, and gaming applications.

Game environments and agent-based frameworks have steadily expanded their models of social interaction between virtual agents. Popular open-world game environments such as Skyrim, Fable, and The Sims use social ties and interactions to drive agent behavior. SOAR and other long-standing agent architectures have more recently been used to model social agents (Li et al., 2008). Newer cognitive agent architectures such as Construct, CLARION, and PMFServ implement off-the-shelf social dynamics (Schreiber & Carley, 2007; Sun, 2007; Silverman, Bharathy, Nye, & Eidelson, 2007).

With agents becoming visually and conversationally realistic, the next frontier of behavioral realism is the interaction between virtual agents. Most commonly,

Presence, *Vol. 22, No. 2, Spring 2013,* 110–140
doi:10.1162/PRES_a_00144

*Correspondence to benjamid@seas.upenn.edu.

multi-agent immersive environments are populated by behaviorally identical archetypes, scripted individuals, or a hybrid of these types. This leads to noticeable repetition and monotony. A longitudinal examination of human–agent interaction by Bickmore, Schulman, and Yin (2010) identified repetitiveness as a primary user complaint in dealing with a virtual agent. The traditional solution to repetition is costly: adding more behaviors for agents.

Worse, more behaviors do not equal more realism. Reliance on static action sets inherently reduces the realism of virtual agents in immersive environments: real societies go through trends with emergent cliques participating in similar behaviors. Expanded action sets alone cannot introduce such trends. As such, adding more behaviors makes the *virtual agents* more real but does nothing to improve the realism of the *virtual agent society*. Rather than adding behaviors, social ties between virtual agents can be used to make behavior more dynamic.

Social learning and adoption of new behavior can be used to represent a more realistic virtual agent society. Learning makes action sets dynamic, allowing new behaviors to supplement and replace old ones. It also allows dynamics such as competing behaviors to emerge. This process can increase realism in game environments and extend social simulation to new problems. Agents in virtual environments lack these capabilities for two primary reasons:

1. Lack of support by agent architectures.
2. Fear of losing control of the agents.

First, agent architectures typically used to drive virtual agents lack key mechanisms supporting the adoption of new behavior. Commercial virtual environments, such as open-world games, often treat interaction between virtual agents as window dressing, rather than as a mechanism driving game state. Social network simulations model adoption in terms of structural factors and use very simple agents, if any (Delre, Jager, Bijmolt, & Janssen, 2010; Centola, 2010). Complex adaptive systems models, such as Rogers, Medina, Rivera, and Wiley (2005), model diffusion patterns due to social factors but utilize higher-level anthropological and sociological mechanisms (e.g., homophilly) rather than lower-level cognitive mechanisms (e.g., attention processes). Cognitive agents capture these lower-level mechanisms, but their application to studying the spread of behavior has been limited. Overall, social network models that utilize social factors have not been extensively applied to immersive environments, except for mechanisms such as flocking or social contagion.

Second, F. Dignum (2012) hypothesizes that developers of virtual environments, such as game designers, are concerned about losing control of the game. Agents learning and adopting new behaviors pose a clear risk in this regard: if an agent might learn any arbitrary action, what would prevent it from acting erratically? Given that established agent architectures have not yet answered this question convincingly, this is a genuine concern. Random or unrealistic adoption trends will hinder immersion in the virtual environment.

To implement the believable spread of behavior in an immersive virtual environment, the question is: "what factors drive social learning by humans?" From the standpoint of an agent, this boils down to *who* it learns from, *what* actions it prefers to learn, and *how* this information reaches the agent. While these mechanisms are not well-explored in virtual agents at the cognitive level, a large body of literature studies the factors that drive social learning in humans (Bandura, 1986). This literature was used to develop a biologically inspired cognitive model for agents. This model emulates the mechanisms that determine who humans learn from and what actions they tend to adopt.

Building on this model, a set of agent-based simulations explores the advantages of this approach. These simulations model two competing actions spreading in Hamariyah, a virtual Iraqi village based on human terrain data provided by the U.S. Marine Corps (Silverman, Pietrocola, et al., 2009). These simulations extend the NonKin village framework (Silverman et al., 2012), using the new cognitive model to drive agents. In this paper, NonKin is used as a simulation environment, to examine adoption patterns. However, the NonKin framework is primarily used to drive agents in immersive training environments, and the results demonstrated here can be directly ported into an interactive real-time

3D environment. These simulations highlight the potential for cognitive agents to enhance the realism and analytical power of agent-based simulations for studying the spread of behavior.

2 Socially Learned Behavior: Prior Work

Prior work on adoption of behavior by agents has not focused on immersive virtual environments, so this review examines a broader range of socially learned behavior by agents. Additionally, this research focuses on descriptive modeling of human social learning, so this literature review only considers systems that model one or more theories of human social cognition. This is an important distinction, since normative agents model optimal behavior (a rational agent), while descriptive agents model human behavior. Existing work falls into two main categories: social-network simulation and teachable agents.

Social-network simulations have recently been applied to model the spread of healthy behavior in online communities, meme utterances, and diffusion of innovations (Centola, 2010; Gruhl, Guha, Liben-Nowell, & Tomkins, 2004; van Eck, Jager, & Leeflang, 2011). Pure social-network simulations lack a complete virtual environment: agent properties, social ties, and update rules comprise the full simulation state. In some social simulations, network topography is the only independent variable, so the models are sometimes agent-based in name only. One notable exception is Construct, a multilayered social-network architecture (Schreiber & Carley, 2007). Since Construct models agent communication of information, new behavior is one type of information that Construct agents can learn. Construct agents are part of the larger class of organizational modeling (see V. Dignum, 2009, for an overview of related approaches). However, as an organizational model, Construct agents focus on high-level group dynamics rather than individual behavior in a virtual environment.

Complex adaptive systems (CASs) also move beyond basic networks, using adaptive agents within social network simulations. As a theoretical concept, complex adaptive systems cover most meaningful agent-based simulations with any degree of adaptation or emergence (Holland, 1998). However, from a literature standpoint, CAS simulations that include social learning typically use lightweight agents and depend on one or two simple mechanisms that implement normative theories, such as game-theoretic agents (Panait & Luke, 2005). Voting mechanisms, social norms, and coordination games have frequently been modeled using these approaches (Lim, Stocker, Barlow, & Larkin, 2011; Van Segbroeck, de Jong, Nowe, Santos, & Lenaerts, 2010). CAS approaches are seldom designed to withstand scrutiny as individual agents; instead, their power lies in their emergent patterns (Railsback, 2001). For an immersive simulation, individual differences between agents are pivotal, because users will interact with them and develop expectations. Cognitive agents more commonly model these aspects and have been used to drive virtual agents (Sun, 2007; Silverman, Bharathy, Johns, et al., 2007; Laird, 2008). However, these agents have not focused extensively on the spread of behavior, as this is typically studied at the societal or organizational level.

Teachable and imitative agents are a second major topic in social learning of behavior (Knox, Fasel, & Stone, 2009). Agents are taught behaviors for two primary reasons: to teach the teacher (teachable agents), or to teach the agent (imitative agents). Unlike social simulations, teachable agents often represent a range of domain behaviors that are taught through dyadic interaction or small societies. For learning environments, teachable agents help the user solidify knowledge and skills through a pedagogy known as learning by teaching. Such agents are increasingly common and have been applied to the instruction of math, language, and metacognitive skills (Pareto et al., 2011; Blair, Schwartz, Biswas, & Leelawong, 2007).

In robotics, learning behavior in a physical environment is a difficult task. To address this challenge, imitative robots have been designed to learn from demonstrations by a human, or by another robot performing the action (Billard & Dautenhahn, 1999). In some cases, imitative robots only learn affordances (opportunities for action), while in other cases, they infer intentionality and model true imitation (Zentall, 2007). Multi-robot teams have also used communication-based imitation

Table 1. *Coverage of Contemporary Agents for Learning New Behavior*

Behavior learned	Level of analysis Dyadic	Micro/meso	Macro
Skills (how to)	Imitative agents Teachable agents	Imitative teams	—
Affordance (what)	Imitative agents	Agent-based simulation Imitative teams	Social simulation
Intentionality (why)	Imitative agents	—	—

to speed up learning of the behavior space (Barrios-Aranibar, Alsina, Nedjah, Coelho, & Mourelle, 2007). In virtual worlds, teachable agents have similarly been taught language and been trained to recognize behaviors (Kerr, Hoversten, Hewlett, Cohen, & Chang, 2007; Kerr, Cohen, & Adams, 2011).

Table 1 summarizes the type of behavior learned by different types of agents, and the level of analysis of the behavior learned. Agents socially learn three distinct but related aspects of behavior: skills (how to do it), affordances (what can be done), and intentionality (why to do it). This research focuses on the center of the table: using agent-based simulation to model individual and group-level social learning of affordances. Socio-cognitive agents are used to model appropriate behavioral interactions between agents (at the micro-level) and the emergent spread of behavior by groups of agents (at the meso-level). These agents are designed to learn new affordances: opportunities for action.

Based on these targets, this research attempts to satisfy three conditions:

1. Realistic agent actions.
2. Social learning about new action opportunities (affordances).
3. Realistic adoption of actions by agent network clusters/groups.

The first condition is well-addressed by current lines of research: many projects exist to make the actions of an individual agent visibly, audibly, and rationally plausible. While major challenges remain for bringing individual agents to the next level, this class of problems has been explored extensively. This research builds on NonKin village, a framework that connects to 3D environments and models patterns of daily life (Silverman et al., 2012). NonKin was developed as a training environment for cultural skills, and it handles action representation and presentation capably. As such, this research focuses on the second and third conditions.

On its own, the second condition is nearly trivial: it is easy to support affordance learning by virtual agents, so long as you don't care about who learns or adopts the new actions. Simple social contagion mechanisms are sufficient to satisfy this condition. However, such mechanisms violate Dignum's constraint, since they "lose control of the game" (F. Dignum, 2012).

The third condition imposes this constraint: adoption patterns must be plausible. Unlike the other challenges, realistic adoption patterns by cognitive agents is a relatively new area. Prior work has not shown that the spread of behavior in an virtual agent society can be modeled such that agents continue to act coherently with their track record of actions. This paper approaches that problem by developing an agent-based cognitive model intended to support realistic social learning and adoption of affordances within an immersive virtual environment.

3 Modeling Affordance Transmission

To model the spread of behavior, this research focuses on socially transmitted affordances. The ecological approach to perception posits that the environment is perceived in terms of the affordances that it offers, referred to as direct perception. Affordances always

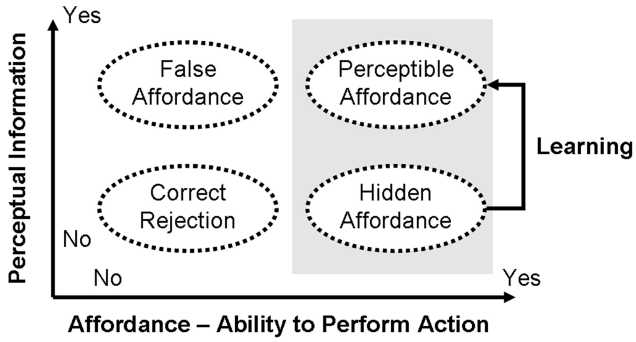

Figure 1. *Relationship between affordances and perception.*

exist: they represent the potential for action (Gibson, 1986). For example, a human has the affordance to swing a hammer. A goldfish does not have this affordance, as it has no hands. Autonomous agents often fit this ecological model: they typically have a static set of capabilities, but may have a wide variety of opportunities for action in their environment. As such, agents are commonly not learning actions in terms of behavioral movements, but are instead discovering affordances: their possibilities for action in an environment.

Affordances are not always known, however. As shown in Figure 1, Gaver (1991) framed this issue using two orthogonal aspects: (1) Is an affordance available? and (2) Is the affordance perceptible? For example, a hidden light switch always offers the affordance to be turned on by pressing it. However, until the switch is identified, it represents a "hidden affordance." A hidden affordance is a potential for action that an organism is not aware of yet. By learning an affordance, an agent moves from having a hidden affordance to having a perceptible affordance (known affordance). In this way, an agent becomes aware of a new action opportunity. Social learning of affordances is important because the space of possible actions can be vast.

3.1 A Memetic View of Affordance Learning

For modeling purposes, socially learned affordances were framed as a type of meme. A meme is a unit of cultural information that spreads by repeated reproduction from one agent to another (Dennett, 1995). A model for meme transmission was synthesized from Bandura's social-learning theory and Shannon's information theory as shown in Figure 2 (Bandura, 1986; Shannon, 1948).

These theories provide complementary processes for examining the flow of information between and within individuals, respectively. The social-cognitive theory establishes the necessary stages for an agent to repeat socially learned behavior: attention to the behavior, retention of the affordance, motivation to repeat the behavior, and physical production of the behavior (Bandura, 1986). However, social-cognitive theory offers little insight for the transmission of information through the environment. Information theory addresses transmission through an environment explicitly, where a source transmits through a medium to a receiver to reach a destination.

This framework offers a comprehensive view of meme transmission in terms of agents sharing a common environment. This framework is particularly well-suited to modeling the spread of socially learned affordances, as the information of an affordance directly corresponds to behavior. Additionally, the separation between Bandura's four cognitive phases of adopting new behavior helps ensure coherent agent behavior. Since learning a new action does not entail motivation to repeat it, an agent learning a new, unattractive behavior would never reproduce it. This has the dual effect of keeping individual behavior realistic, while also slowing the diffusion of that behavior to that agent's social ties.

Notably, this framework does not explicitly address directed communication: agents telling each other about an affordance. This is by intention: agent communication *is* a behavior. Additionally, the mirror neuron hypothesis posits that language emerged from observational learning (Arbib, 2011). As such, verbal communication may best be viewed as a second-order process for transmitting affordance information. Many agent-based frameworks represent communication as a separate process that is not subject to the same restrictions as a standard action (Schreiber & Carley, 2007; Panait & Luke, 2005). This framework takes the opposite view: communication is a behavior that must compete with the agent's other opportunities.

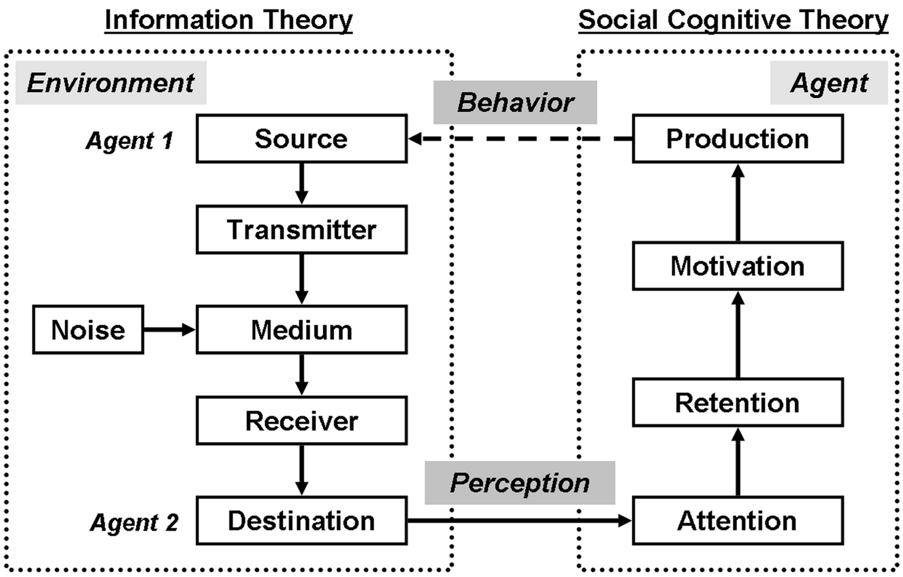

Figure 2. *Systems model for meme transmissions.*

By implication, this means that communication may be ignored by its intended recipient or observed by unintended recipients. Framing communication as a behavior allows the agent's environment to determine its affordances for communication opportunities (e.g., who they can talk to, the mediums available, etc.). While this paper focuses on observational learning, the framework naturally extends to communication as well.

3.2 Cognitive-Agent Architecture

Based on this systems model for affordance transmission, a cognitive model was created using the PMFServ socio-cognitive architecture. PMFServ implements cognition using a model-of-models approach: integrating best-of-breed social science models and performance moderator functions (PMFs) to form a cognitive model (Silverman et al., 2012). These models incorporate the OCC cognitive structure of emotions (Ortony, Clore, & Collins, 1988), GLOBE cultural traits (House, Hanges, Javidan, Dorfman, & Gupta, 2004; Hofstede, 2003), Hermann's leadership traits (Hermann, 2005), affordance-based perception (Gibson, 1986), subjective utility (Damasio, 1994), and multiple other well-supported moderators of cognition and decision-making. While reviewing its existing features in detail is beyond the scope of this paper, PMFServ has a long track record for modeling decision-making and has been used to drive agents in crowd environments (Silverman, Johns, Cornwell, & O'Brien, 2006), leader decision games (Silverman & Bharathy, 2005), and country stability simulations that had an accuracy of over 85% (Bharathy & Silverman, 2010; O'Brien, 2010).

An attractive feature of the PMFServ framework is that agents employ affordance-based perception (Silverman et al., 2006). However, PMFServ's standard agent perceives all of the affordances of its environment and lacks any cognitive mechanisms for managing attention and retention of new affordances. To simulate affordance transmission, significant additions to the PMFServ model base were required. The following section discusses the theories implemented as models, how these theories interact with existing PMFServ models, and how these models help to model social learning and adoption of new behavior.

3.3 Attention Mechanism

Attention is a fundamental mechanism for social learning and the spread of new behavior. Without

Algorithm 1 Attention algorithm

$E_{\text{Att}} = \{\ \}$
for $i = 1$ to N **do**
  ATTENDED_EVENT = $X(E, E_{\text{Att}})$
  if ATTENDED_EVENT != No Event Attended **then**
    $E_{\text{Att}} = E_{\text{Att}} \cup \{$ ATTENDED_EVENT$\}$
  end if
end for

attention, a cognitive agent cannot demonstrate the cocktail-party effect and other cases where an agent differentially processes some stimuli over others (Cherry, 1953). In social network models, attention is often represented as relatively random. However, a multitude of findings demonstrate that the cognitive mechanisms for attention to events are far from uniformly random. As such, attention was driven by a mixture of cues that will be described in the following section.

This attention model corresponds to a series of winner-take-all competitions for attention between simultaneous events, a process that has some support in neurological research (Lee, Itti, Koch, & Braun, 1999). Attentional salience determines the probability that an agent will attend to an event. This is accomplished by first calculating a salience for each event occurring during a time step. An additional salience term exists to represent inattention salience: the salience of background events not simulated that might be attended to instead of the simulated events. This vector of saliences is normalized to form a probability vector, from which a finite number of events are chosen. Each event is chosen without replacement, except for inattention, which always remains an option.

The algorithm for drawing the set of attended events is displayed as Algorithm 1, where N is the maximum simultaneous events attended, E is the set of all current observable events, E_{Att} is the set of currently attended events, and $X(E, E_{\text{Att}})$ is a random variable returning at most one unattended event from the set $E \backslash E_{\text{Att}}$. The output of this algorithm is E_{Att}, the total set of attended events. If $X(E, E_{\text{Att}})$ returns no event, this represents inattention and one less total event will be attended. This attention algorithm is effectively an iterated drawing from the yet-unattended events, with some probability of no event being attended. Attended events are processed by the learning model, which can learn new affordances.

$$P(e, E, E_{\text{Att}}) = \begin{cases} \frac{s_e}{s_I + \sum_{e \in E \backslash E_{\text{Att}}} s_e} & \text{if } e \in (E \setminus E_{\text{Att}}) \\ \frac{s_I}{s_I + \sum_{e \in E} s_e} & \text{if no event attended} \\ 0 & \text{if } e \in E_{\text{Att}} \end{cases} \tag{1}$$

The probability that an event (e) receives enough attention to be processed cognitively is determined by the distribution of $X(E, E_{\text{Att}})$ and will be referred to as $P(e, E, E_{\text{Att}})$. The probability distribution for choosing an event to attend is shown in Equation 1, where E is the set of all simultaneously observable events, E_{Att} is the set of events already attended to, s_e is the salience of an individual event e, and s_I is the inattention salience. Events with higher salience are more likely to be selected, as they fill a greater fraction of the probability vector. However, for attention to work realistically, it must be based on appropriate cues from cognitive and social psychology.

3.4 Attention Cues

Attentional salience is calculated as a function of attentional cues. Any action involves an actor (source), behavior (action), and some outcomes (results). Theories of attention and persuasion both indicate that attentional salience is influenced by central and peripheral cues (Treisman & Gelade, 1980; Petty & Cacioppo, 1986). Figure 3 displays how an observing agent breaks an event down into a set of cues that are used to determine attentional salience. Due to space limitations, each cue will only be described at a high level, but for the interested reader, further theoretical and technical details on their implementation are contained in Nye (2011).

For an affordance, central information includes direct information about the associated behavior. These include whether an agent can perform the observed action, whether the action resulted in appealing outcomes, or whether the action seems new. These

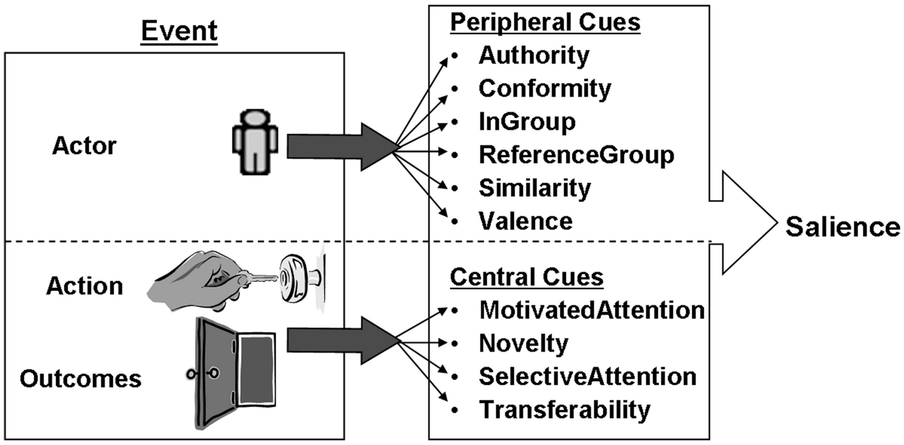

Figure 3. *Event attention cues.*

influences are known as transferability (Bandura, 1986), motivated attention (Fazio, Roskos-Ewoldsen, & Powell, 1994), and novelty (James, 1890), respectively. Selective attention was also used as a cue, so that agents could choose to pay more attention to a particular agent (Simons & Chabris, 1999).

Peripheral cues such as social-factors cues are equally important for directing attention, however. Social influence is commonly implemented in social networks, but is often represented as a single intrinsic agent property. The problem with this approach is that social influence is a multi-faceted, relational construct. To address this issue, social influence was represented by implementing multiple established theories of social influence.

The social cues implemented were authority (Mantell, 1971), conformity (Tanford & Penrod, 1984), similarity (Platow et al., 2005), valence (Hilmert, Kulik, & Christenfeld, 2006), in-groups (Tajfel, 1982), and reference groups (Kameda, Ohtsubo, & Takezawa, 1997). These cues represent some of the most well-established factors of social influence. Authority influence is the additional influence due to an actor's leadership or authority positions. Conformity is the added impact of observing multiple actors performing the same behavior. Valence is the amount that an observer likes the actor performing a behavior. In-group influence is the additional weight given to a member of the same group. Reference group influence is the additional weight based on membership in a group that an agent uses as a comparison (e.g., keeping up with the Joneses). The following sections discuss each of these factors in further detail.

The total salience of each event is calculated using a linear weighted sum of these cues (i.e., $s_e = w_1 \times \text{Authority} + w_2 \times \text{Conformity} + \cdots + w_{10} \times \text{Transferability}$). Since the relative strengths of these factors are not well-studied, best guess weights were calculated from their observed effect on either attention, perception, or retention. A linear sum was chosen based on the KISS principle, as it was the simplest way to combine cues into a total salience (Axelrod, 1997). While there are good reasons to believe that some of these factors interact, psychology literature has not yet produced the studies that demonstrate how these factors interact.

3.4.1 Novelty (Central). The three central cues modeled were novelty, motivated attention to outcomes, and transferability. Novelty indicates how new a stimulus appears (James, 1890). Novelty decreases with respect to the number of prior exposures stored (Johnston, Hawley, Plewe, Elliott, & DeWitt, 1990). To model this, novelty is calculated as a function of an agent's familiarity with each action and agent present at an event. The novelty model calculates this based on familiarity levels from the memory model, which will be described in Section 3.5.

For any given event, the novelty is calculated as the root mean square of the familiarity values of the actor of the event and the action of the event. The novelty calculation for an event is shown in Equation 2, where f_{Actor} is the familiarity of the event's actor, and f_{Action} is the familiarity of the event's action according to the memory model.

$$\text{Novelty(Event)} = \sqrt{0.5((1 - f_{\text{Actor}})^2 + (1 - f_{\text{Action}})^2)} \quad (2)$$

This representation was chosen because it allows a high degree of novelty if either component is novel. This dynamic was chosen because it allows representation of processes such as dishabituation, where adding an additional stimulus can restore the response to a habituated (familiar) stimulus. In this context, the response of interest is active attention. This implementation allows a return to novelty when a highly familiar person suddenly engages in a totally new action. Conversely, if a straight average was used, then a completely familiar person could be at most 50% novel. Alternately, taking the maximum novelty component would give no extra credit for a new person taking a new action. A root mean square parsimoniously represents these important dynamics within the simulation.

3.4.2 Motivated Attention (Central). Motivated attention refers to the tendency of humans to pay more attention to objects or events that are relevant to their goals or needs (Fazio et al., 1994). For example, a hungry person is more likely to notice someone eating. Motivational cues are handled by allowing agents to analyze the outcomes of events that occur in their presence.

PMFServ's core cognitive models evaluate their potential actions based upon activations that determine the attractiveness of that action, as mediated by their values and beliefs (Silverman et al., 2006). These mechanisms for motivation will be discussed in Section 3.6. To calculate a factor for motivated attention, an agent processes an event that results from some other agent's action. In processing this event, the agent calculates the subjective expected utility (SEU) as if the agent had been the actor in that event and the outcomes were the same. So, for example, if agent B is eating a sandwich, the motivational salience for agent A is a function of the subjective benefit for agent A eating that sandwich (even if no more sandwiches currently exist to eat).

$$\text{MotivatedAttention(Event)} = 0.5 \times (1 + \operatorname{sgn}(\text{SEU}_{\text{Event}}) \times (|\text{SEU}_{\text{Event}}|^{0.25})) \quad (3)$$

Equation 3 displays the central motivated attention calculation for an agent observing a given event (Note: the sgn symbol represents the sign function, producing -1 for negative values and 1 otherwise). $\text{SEU}_{\text{Event}}$ represents the SEU of activations that the perceiving agent would receive had it been the actor in that event and the outcomes were the same. An adjustment to the raw utility rescales the value from the utility's range of $[-1,1]$ to $[0,1]$. The second rescaling factor takes the fourth root of the absolute SEU value. This factor was introduced during model calibration to adjust the small range over which SEU typically operates in PMFServ (about $[-0.05, 0.05]$) to cover a motivation range closer to $[0.25, 0.75]$.

3.4.3 Transferability (Central). The third central cue modeled was transferability. Transferability influence refers to the additional influence conferred by an agent having similar capabilities and doing actions that could be imitated. Often, this trait is studied in children at different developmental stages. Children have a preference to attend to and imitate those of similar ability level on tasks (Bandura, 1986).

The transferability influence model allows agents to process an observed event and determine whether they could do the same action at the current time. This determination is only based upon the agent's current affordances at the particular moment, not on any past or potential affordances. This implementation has the advantage of easily classifying events into those that they could imitate (transferability = 1) and those that they could not (transferability = 0).

3.4.4 Authority (Peripheral). Six peripheral cues were also incorporated into the model, representing social cues. The authority influence model represents the additional influence conferred by a position of authority. The effects of authority on behavior have been well documented by Milgram (1963) and by Mantell (1971). PMFServ represents the authority of agents within their respective groups (Silverman et al., 2006). Since this factor is already represented, the authority submodel wraps this factor for use as a social cue.

3.4.5 Conformity (Peripheral). The conformity model has its theoretical roots in the seminal work done by Asch (1955). Later work by Tanford and Penrod (1984) proposed the social information model (SIM), a probabilistic conformity influence function. Their analysis produced a curve as stated in Equation 4, where S is the number of conforming sources, and T is the total number of non-conforming targets.

$$\text{ConformityInfluence}(S, T) = e^{-4\times e^{\frac{-S^{1.75}}{T}}} \tag{4}$$

The implemented conformity model uses this equation verbatim. However, the context of its usage is slightly different from that of the original SIM model. While that model assumed a set of confederates, these models assume that agents act based upon their own opinions but still exert influence. As such, any set of agents engaged in a particular activity forms a group of influence sources (S). The remaining agents involved in other activities are the target group (T). As such, agents can calculate the conformity influence of any activity in the simulation for any given action occurring at the time.

3.4.6 Similarity (Peripheral). The similarity model calculates a social-influence factor based upon how much an agent feels it has in common with another agent. The influence of similarity on attention and influence has been an important topic in the domains of social psychology and social-network analysis (Platow et al., 2005). PMFServ contains a model that estimates a proxy for similarity, known as Goals, Standards, and Preferences (GSP) congruence (Silverman et al., 2006). This estimate is based on the GSP personality model, which is described in Section 3.6.1. The GSP model expresses an agent's personality as a tree of traits connected by weighted links. Each weight determines the importance of a child trait toward a parent trait (e.g., 40% of an agent's goals focus on safety). GSP congruence is calculated by transforming agents' GSP trees into vectors of normalized linear weights and calculating the nearness between these vectors. The standard GSP congruence function is shown in Equation 5, where $\vec{w}$ is the perceiving agent's GSP vector, $\vec{w^*}$ is the observed agent's GSP vector, and N is the number of elements in $\vec{w}$.

$$\text{GSPCongruence}(\vec{w}, \vec{w^*}) = 1 - \frac{\sum_{i=1}^{N}(\vec{w_i} - \vec{w_i^*})^2}{\sum_{i=1}^{N}(\vec{w_i})^2 + (\vec{w_i^*})^2} \tag{5}$$

The similarity influence model builds off of the GSP congruence model, using GSP congruence as a similarity term. By allowing agents to detect this factor without noise, the model assumes that the agents generally estimate an accurate perception of similarity. This model is best applied when agents have prior knowledge about other agents' personalities or quickly assess other agents' personalities. Even where agents are not familiar, it provides a useful first-order estimate of the perceived similarity.

3.4.7 Valence (Peripheral). Valence influence is caused by general like or dislike of another person or group. This is related to the halo effect, whereby an attractive person appears more competent (Kelley, 1955). Works such as Hilmert et al. (2006) have experimentally shown that valence affects social influence. PMFServ valences are directed properties of one agent toward another entity. Valence influence exposes these properties as cues for attention. Since valence ranges from $[-1,1]$ in PMFServ and all cues are fitted into a range of $[0,1]$, a small transform is applied to valence values to rescale and shift them into the appropriate range.

3.4.8 In-Group (Peripheral). The in-group influence model represents the social influence based on membership in a mutual group or clique (Tajfel, 1982). PMFServ has a structure for representing group membership, which allows members to be part of a group. This cue determines whether agents share a common group (in-group = 1) or share no common groups (in-group = 0).

3.4.9 Reference Group (Peripheral).

Reference-group influence represents the influence based on an agent belonging to a group against which an agent self-compares, such as a desirable group (Kameda et al., 1997). PMFServ has an analogous factor within its model set that is an agent's internal membership with a group (Silverman et al., 2006). Internal membership measures how much an agent desires to participate and support a group. As this measure is explained by Eidelson (2007), it will not be covered in detail here.

Reference group influence uses a variant of PMFServ internal membership that has been scaled to fit into a range of [0,1]. This model can report back the desire to belong in any given agent's group (if belonging to a group). This value can be independent of in-group influence, since people are not always a member of their preferred group.

3.4.10 Selective Attention.

Selective attention is a construct that refers to the additional probability of perceiving events performed on an object that an agent actively perceives, as opposed to other peripheral events (Simons & Chabris, 1999). Selective attention is implemented by having agents keep a record of the objects and agents they are actively attending to at the current time. PMFServ agents are able to actively take actions on other agents, including actions of active perception (watching).

$$\text{SelectiveAttention}(x) = \begin{cases} \frac{1}{N} & \text{if } x \in X_{\text{Targeted}} \\ 0 & \text{if } x \notin X_{\text{Targeted}} \end{cases} \quad (6)$$

As such, the selective-attention model records all entities that an agent is currently engaged in action upon. This means that selective attention is focused on any targets being watched or acted upon by an agent. Selective attention is spread uniformly across these targets as noted in Equation 6. This allows agents to choose who will be the target of their selective attention, as is observed in the cocktail-party effect (Cherry, 1953).

3.5 Retention Mechanisms

Since this cognitive model was primarily intended to address the issue of who it is that learns and adopts new affordances, the memory model was kept as simple as possible. Many affordances of interest are relatively simple and memory effects are not the main barrier to adoption. As such, memory was implemented as a simple associative structure. Associative memory works by strengthening connections between elements, stimuli, or constructs due to repeated pairing (Mackintosh, 1983).

This information is used for two purposes. First, this memory model supports affordance learning. Once an action is stored in the agent's memory, the affordance for that action becomes known. As such, attending to an event with a new behavior will let the agent learn this behavior. Second, the model is used to calculate familiarity because this is needed to determine the novelty of events.

$$\text{Familiarity}(\text{Entity}) = 1 - e^{-r_f \times N_E} \quad (7)$$

The familiarity equation is stated in Equation 7. The input to the equation, Entity, is an action, agent, or other entity contained within a learned pattern. N_E is the number of exposures to that entity and r_f is a familiarity rate that determines the steepness of the curve. Within the current implementation, r_f was set to 0.2, as this allows familiarity to reach 95% after 15 exposures. Empirical research indicates that the exposure effect hits its maximum after between 10 and 20 exposures, so this seemed to be a reasonable familiarity rate (Bornstein, 1989).

3.6 Motivation Mechanisms

The motivation to perform an action is controlled by PMFServ's decision model. As PMFServ's decision model has undergone over 10 years of development, fully understanding these processes requires careful reading of a number of prior papers (Silverman et al., 2006; Silverman, Bharathy, Johns, et al., 2007; Silverman, Bharathy, Nye, et al., 2007). PMFServ agent motivation is driven by a process known as cognitive appraisal theory. In cognitive appraisal, an actor has a set of three models: a GSP personality model, an emotion model, and an SEU decision model. While the models that drive motivation were not modified for this research, they are used by the motivational-attention model, and also determine agent adoption decisions. Due to their importance in determining what actions are expressed, the mechanisms of motivation will be discussed in this section.

3.6.1 Goals, Standards, and Preferences Model. From the standpoint of affordance adoption, the most important model is the GSP model, which stores an agent's personality. Specifically, a personality is modeled by a tree of weights (Bayesian prior odds) that represent the relative importance of each GSP node (trait) to that person. The GSP nodes used in the experiments in Section 4 are listed in Table 2. The nodes in this tree are based on Maslow's Hierarchy (Maslow, 1943), cultural dimensions of organizations (Hofstede, 2003; House et al., 2004), and Hermann's trait analysis of leadership styles (Hermann, 2005). Nodes are split into three main branches: goals (short-term goals), standards (how to accomplish things), and preferences (long-term wishes).

GSP tree factors are based on trait theories, which posit that personality traits are relatively stable over time. As such, this model captures individual differences between agents and determines the outcomes they prefer. However, even agents with the same GSP tree often display very different behavior due to different experiences (e.g., observed events, emotional states), knowledge (e.g., affordances, familiarity), and external contexts (e.g., different roles, economic condition, location). This results in path-dependent behavior, particularly in a multi-agent system. For example, if two initially identical agents compete in a race, one agent will experience winning and the other will experience losing. The agents' behavior will diverge due to the different experiences and any external changes (e.g., prizes, rewards, changes in perception by other agents).

Agents evaluate their experiences in terms of activations. Activations are part of the outcomes of actions that are afforded to agents. Each activation positively or negatively targets a single GSP node. For example, gaining money creates positive activations for a materialistic preference. An action that results in pain for the agent will give negative activations for a safety goal. Similar to attention, social models also impact motivation. As noted earlier, agents have models for group membership and valence (like/dislike) toward other agents and groups. Actions that affect an agent's in-group will activate nodes such as own_people and for_the_group. Similarly, an agent's valence toward an agent or group influences the activations on the be_relationship_focused node produced by action outcomes (e.g., hostile actions toward friends create negative activations on this node).

3.6.2 OCC Emotion Model. The emotion model calculates a set of emotions based on the activations to different parts of the GSP tree (Silverman et al., 2006). The emotions calculated by this model are based on the OCC (Ortony et al., 1988) formalization of emotions. Joy, distress, pride, shame, liking, disliking, gratification, and remorse are emotions considered by the decision model. Pairs of positive and negative emotions are determined by positive or activations to each branch of the GSP tree: goals (joy/distress), standards (pride/shame), and preferences (liking/disliking). Each of these emotions is a normalized vector projection of activation values onto their corresponding node weights (e.g., $\overrightarrow{\text{Weights}} \cdot \overrightarrow{\text{Activations}}$). Gratification and remorse are compound emotions based on the other emotions. Emotions accumulate as a result of events and decay over time. For example, as goals are satisfied, the agent will receive less emotional impact from them, allowing the agent to focus on other goals.

3.6.3 Decision Model. The decision model calculates an agent's SEU for each afforded action based on these emotions (Silverman et al., 2006). The SEU of an action is determined by calculating the expected change in emotions from the activations of an action. Equation 8 displays how the decision model calculates a subjective utility based on the emotions. The expected utility is otherwise calculated in the typical way, based on the probability of action outcomes with different activation sets.

Table 2. *GSP Personality Factors*

Node Name	Description of trait
GOALS: Short-term goals, which connect to joy and distress	
Individual	Overall individual goals, e.g., Maslow (1943) Hierarchy
Belonging	Social acceptance and feeling situated among peers
Esteem	Feeling of self-efficacy and respect
Physiology	Basic bodily needs, such as eating and sleeping
Safety	Personal safety and well-being
STANDARDS: Standards for behavior, how an agent prefers to accomplish tasks. Connects to pride and shame.	
Conformity_assertiveness	Overall importance of conformity and individuality
Assert_individuality	Expressing individuality
Conform_to_society	Conforming to culture
Respect_authority	Showing respect for authority figures
Exercise_of_power_and_culture	Overall importance of power balances in actions
Be_controlling	Controlling others by using power
Be_open	Being open to others, allowing freedom
Honesty	Overall importance of honesty and dishonesty
Keep_one's_word	Keeping promises, being honest
Use_duplicity	Lying for its own sake
Humanitarian_sensitivity	Overall importance of considering lives and showing respect for life
Respect_for_life	Respecting and being sensitive to the lives of others
Disregard_for_life	Disregarding and being insensitive to others' lives
Military_doctrine	Overall importance of adhering to military codes
Shun_violence	Avoiding violence
Use_asymmetric_attacks	Attacking unevenly, even unfairly
Use_conventional_attacks	Use of force-on-force conventional tactics
Scope_of_doing_good	Overall importance of doing good for others
Bring_about_greater_good	Doing good in the world, in general
Look_after_narrower_interests	Only looking after one's own interests
Task_relationship_balance	Balancing tasks and relationships
Be_task_focused	Concentrating on tasks only
Be_relationship_focused	Building relationships or social networks
Treatment_of_out-groups	Overall importance of interaction with out-groups
Out-groups_are_legitimate_targets	Targeting out-groups for discrimination
Enemy_is_out-group	Targeting one's enemies negatively
Friend_is_out-group	Targeting one's friends negatively
Neutral_is_out-group	Targeting neutral parties negatively
Treat_with_fairness	Treating everyone equally

(continued)

Table 2. *(Continued)*

Node Name	Description of trait
PREFERENCES: Long-term wishes for the world state. Connect to like and dislike emotions.	
Desirable future	Actions that produce good outcomes, by scope
For_everybody	Benefit for everyone in society
For_the_group	Benefit to one's immediate in-group(s)
For_the_self	Benefit for one's self
People	Long-term outcomes for specific people, by relationship
Enemy_faction	Long-term outcomes for enemy factions
Friendly_faction	Long-term outcomes for friendly factions
Own_people	Long-term outcomes for own group
Other_groups	Long-term outcomes for neutral groups
Places_and_things	Actions impacting objects or states of the world
Materialistic	Property and monetary objects
Symbolistic	Symbolic outcomes, principles being maintained
Wholistic_spiritualistic	Religious or spiritual matters

$$\begin{aligned}\text{Utility(Emotions)} = \frac{1}{4}((&\text{Joy} - \text{Distress}) \\ &+ (\text{Pride} - \text{Shame}) \\ &+ (\text{Liking} - \text{Disliking}) \\ &+ (\text{Gratification} - \text{Remorse})) \quad (8)\end{aligned}$$

Agents decide on their actions by selecting the option with the highest SEU in the simulations described in Section 4. This means that actions compete against each other to be an agent's top decision choice. Since an agent's emotions depend on their GSP model, agents with different GSP weights tend toward different types of behavior. Finally, agents can only consider actions that they perceive as affordances, so the new attention and memory models also influence action choices. Until an agent learns about an affordance, the agent cannot calculate its utility or choose that action.

3.7 Production Mechanisms

Production mechanisms in PMFServ are represented by the actions associated with affordances. These actions depend on the specific scenario and generate observable events when they occur. The ability to perform an action requires a valid affordance for that action in the environment. As such, the ability to produce an action is atomic—an agent is either able or unable to perform an action. As noted earlier, agents are unable to perform an action unless they are aware of its affordance. This makes intuitive sense, as an agent cannot initiate an action without recognizing the possibility of performing that action (i.e., the affordance).

4 Hamariyah Iraqi Village Simulation

Agents using this cognitive model were used to populate Hamariyah, an archetypal Iraqi village based on a human terrain data set. The scenario examines the spread of adoption of two competing behaviors: giving information to the U.S.-backed government and volunteering to plant an IED near a government building. Since this framework had preexisting actions, the spreading behaviors competed against each other and against the existing action set, which primarily models daily life activities. These simulations were performed to examine whether the cognitive agents could fulfill the three requirements listed at the end of Section 2: realistic

Figure 4. *NonKin 3D environment screenshot.*

agent actions, social learning of new actions, and realistic adoption clusters.

This scenario was generated using the NonKin village framework and data provided by the U.S. Marine Corps (Silverman, Pietrocola, et al., 2009). NonKin village is a virtual village engine based on PMFServ agents (Silverman et al., 2012). The NonKin village can drive agent behavior in a 3D environment or run faster-than-real-time as a simulation without graphical support. Figure 4 shows a screenshot of agents congregating in the NonKin immersive environment. This archetypal village was intended to be representative of a village in Iraq. The human terrain data set includes agents' names, familial ties, group memberships, group roles, special skills, key personality traits, land ownership, and employment. Groups in the region are also described, with a focus on their valence relationships (like/dislike) and historical backstory. Given the scope of the NonKin Village project, it is infeasible to explore every aspect in detail. Instead, the following sections will highlight the key scenario features and extensions that were necessary to study competing behaviors within the village. For more detailed information about the scenario, Silverman, Pietrocola, et al. discuss the human terrain data for Hamariyah and Silverman et al. discuss the architecture and advanced features. However, the work described in Silverman, Pietrocola, et al. used an older version of the NonKin village architecture, so Hamariyah was regenerated from the original Marine Corps data.

4.1 Hamariyah Scenario

The Hamariyah scenario contains 72 agents from the Marine Corps human terrain data. These data were used to determine the initial values for all simulations, which will be described here. These agents belong to three distinct ethnic groups: 11 Heremat members, 38 Shumar members, and 23 Yousif members. As group ties are established by ethnicity, these memberships are static. In addition to agents being members of groups, structures in the NonKin village are tagged by their group affiliation. This allows agents to see whether buildings belong to their group, a group they like, or an unfriendly group. These relationships are determined by the group-to-group valences, whose starting values are shown in Figure 5. Agents can also be employed at a job or can be unemployed. The Heremat group is generally friendly to the U.S., and controls the local police force, but is not a very big group. The Shumar ethnic group is primarily Sunni and unfriendly toward all other groups, especially the U.S. group. It is the largest group, with a majority of its members working as merchants or tradesmen. The Heremat and Shumar ethnic groups both have members working as part of the local government. The Yousif ethnic group is a primarily Shia group, with higher than 60% unemployment and religious leaders in higher positions of authority. Employment and group valences may change due to simulation events (e.g., group-to-group attacks, shops closing, etc.).

Hamariyah contains over 50 standard actions that can be taken by agents, on a variety of targets. The availability and attractiveness of opportunities depends on the context (e.g., location, role, nearby agents), current internal state (e.g., emotions, hunger, etc.), and their current information (e.g., known affordances, familiarity). Nothing is scripted, and agents choose actions autonomously. These actions range from complex multistage actions (i.e., go to market and buy food) down to niche actions for forcing entry into a building. The original set of actions was not modified, as it provided the contextual backdrop for examining the spread of behavior. The most common actions that agents take within Hamariyah village are those related to daily life. These actions include moving from one building to another,

From To	Shumar	Heremat	Yousif	US_Group
Shumar	1.000	-0.2000	-0.4000	-0.6000
Heremat	-0.2000	1.000	-0.2000	0.4000
Yousif	-0.4000	0	1.000	-0.6000
US_Group	0	0.8000	-0.2000	1.000

Figure 5. *Hamariyah group valence starting values.*

Table 3. *Hypothesis Condition: Innovator Agent Demographics*

Demographics	GiveInformation	PlantIED
Number of agents in Shumar ethnic group	1 agent	2 agents
Number of agents in Heremat ethnic group	4 agents	0 agents
Number of agents in Yousif ethnic group	1 agent	4 agents
Number of agents employed	4 agents	3 agents
Average valence toward U.S. (in $[-1,1]$)	0.07	-0.6
Authority (in $[0,1]$)	0.17	0.0

entering/exiting buildings, buying food, working, socializing, praying, and sleeping. Agents are also able to take less common actions, such as attacks, shootings, and hiring/firing employees.

4.2 Socially Learned Behavior (Memes)

To examine socially learned behavior, two new behaviors were added to the Hamariyah Iraqi village: give information and plant IED. Both of these behaviors could only be performed on the government meme target structure. The giveInformation action represents acting as an informant to the U.S. Give Information is the learned affordance that an agent can go to the U.S. structure to inform on dangerous members in the village. PlantIED is an opposite and competing action. This action involves volunteering to help plant an IED in the vicinity of the U.S. structure (just volunteering, not actually performing an action to emplace an IED). Both actions have inherent risks that give negative activations for personal safety.

Both simulations were run to convergence, a length beyond where full learning was typically observed. This modeling choice means that the experiments underestimate the number of holdouts, where holdouts are the agents that would never learn or express the action. However, this allows for better examination of relative expression rates and diffusion. By making agents more likely to learn and perform the action at some point in the scenario, the differences between late adopters can be identified instead of clumping into a large class of agents who never perform the action.

4.3 Iraqi Village Experimental Cases

The Hamariyah scenario was run under two experimental conditions: a hypothesis case and a randomized case. These conditions were used to examine differences in patterns between a carefully selected set of innovators and patterns from random sets of innovators. The hypothesis case assumed that a particular set of six agents initially knew each behavior, based upon the agents' roles in society. Table 3 shows some basic demographic information about the agents in the hypothesis condition. In the hypothesis case, the giveInformation behavior was initially known by six agents chosen because they were members of the local police or involved with the local government. Agents in the police force and government could be expected to be aware of how and where to provide intelligence to the U.S. forces in their area. GiveInformation innovators are primarily Heremat and slightly like the U.S. group. The plantIED behavior was initially known by six agents categorized as anti-U.S. and their kinetic special skills listed them as an IED maker or an IED emplacer. The plantIED innovators are primarily tradesmen or unemployed, and they greatly dislike the U.S. This scenario was intended to represent the transmission of competing behaviors under realistic conditions.

The randomized condition started with a random set of agents aware of each affordance, so there was less initial predisposition to spread the behavior, but it might reach a wider variety of agents. At the start of each run, six agents were randomly chosen to start with the giveInformation affordance and another six agents were randomly chosen to start with the plantIED affordance. This condition was intended to examine the patterns of behavioral transmission that exist when actions are available to agents who would not normally be expected to carry them.

4.4 Example of Affordance Transmission

The cognitive model explained in Section 3.2 determines how agents spread these behaviors in the Hamariyah village. To help ground this process, this section offers a small example of an agent socially learning and adopting the plantIED behavior. Assume three agents: a Shumar Baathist militant, a Shumar Al Qaeda Iraqi (AQI) insurgent, and a Heremat shopkeeper. Initially, only the Baathist is aware of the plantIED affordance.

4.4.1 Attention Example. The Baathist performs a plantIED volunteering behavior where both the insurgent and shopkeeper might observe this action. The attention model for each of these observers breaks down this event into attention cues, as shown in Table 4. This table demonstrates that the AQI agent has many more cues that would lead this agent to pay attention to the Baathist's action. Motivated attention, valence; in-group membership, and reference group cues are all fairly high for the AQI agent, but low for the shopkeeper. The motivated attention is higher, because the AQI agent would also be interested in taking a plantIED action. The valence is high, because those two specific agents were designated as friendly during the design of the village; in-group and reference group values are high, because the Baathist and AQI agent are both part of the Shumar ethnic group.

Table 4. *Attention Cues for Observers of Baathist PlantIED*

	Heremat shopkeeper		AQI insurgent
Novelty	1.0	=	1.0
Motivated attention	0.34	$<$	0.64
Transferability	0.0	=	0.0
Selective attention	0.0	=	0.0
Authority	0.0	=	0.0
Conformity	0.19	=	0.19
Similarity	0.61	$\approx$	0.62
Valence	0.0	$<$	0.5
In-group	0.0	$<$	1.0
Reference group	0.35	$<$	0.46

Novelty, transferability, selective attention, authority, conformity, and similarity were fairly similar for both agents. This is because neither observer knows the affordance (high novelty), neither agent is currently in a position to perform the action (not transferable), and neither agent is actively paying attention to the Baathist agent (no selective attention focus). Authority is zero for both because the Baathist does not have authority in any group. Conformity is low because the Baathist agent is the only one performing plantIED out of the three agents (e.g., $S = 1, T = 2$ for Equation 4). Finally, similarity is comparable because both observers have personalities that are equally different from the Baathist.

These cues determine the salience for the attention model. The attention model probabilistically determines whether each observer pays attention to the Baathist's action. These cue sets mean that the AQI agent is approximately twice as likely to attend to the plantIED action. As such, for further discussion, it is assumed that the AQI agent paid attention to the action, but the Heremat shopkeeper did not.

4.4.2 Retention Example. Since the AQI agent attended to the plantIED action, this agent learns from this event through its memory model. Two changes occur for the agent. First, the AQI agent can now perceive the plantIED affordance. The agent permanently learns this knowledge and will select plantIED at any time where the action is afforded and is the action choice with the highest SEU. Second, the agent becomes more familiar with the plantIED action and has a lower novelty toward that action (from 1.0 to about 0.9).

4.4.3 Motivation Example. The AQI agent is then able to select the plantIED action during the decision process, since that agent is now aware of the affordance. When the AQI agent evaluates the action, its activations depend on the expected outcomes. PlantIED means volunteering for a violent asymmetric attack, so the action provides success activations for disregard_for_life, use_asymmetric_attacks, and be_task_focused. Similarly, it has negative activations on safety, respect_for_life, and shun_violence. Since the AQI agent has a low valence toward the U.S. group, plantIED also generates success activations on for_the_group and enemy_is_out-group.

The AQI insurgent's GSP model strongly matches these activations. Its personality gives a high weight to violent traits (disregard_for_life, use_asymmetric_attacks, enemy_is_out-group) and low weight to nonviolent traits (respect_for_life, shun_violence). The AQI agent also places a very low weight on safety goals, so is willing to engage in high-risk actions. This means that the AQI agent should be highly motivated to select the plantIED action. The results discussed in Section 5.3.2 confirm this expectation, as AQI agents were among the earliest adopters of plantIED.

For comparison, the Heremat shopkeeper would be a poor match for the plantIED activations. As the shopkeeper has a positive valence toward the U.S. group, this agent does not receive any activation for nodes such as enemy_is_out-group. The GSP for the shopkeeper also has a high weight for nodes such as safety, respect_for_life, and shun_violence. As such, the shopkeeper would have a negative subjective utility for the plantIED action, and this agent would not generally perform this action.

However, it should be noted that all agents must perform an action regardless of how bad their options are. So then, if the shopkeeper's only action choices were to volunteer to plant an IED or to suffer some cruel fate with worse activations, plantIED could still be selected. As such, expressing an action depends not only on its activations but also on the activations of other actions available. This also means that seemingly unrelated GSP nodes can prevent selecting the plantIED action by leading an agent to prefer other actions, even if the agent has a positive utility for plantIED. This means that the converse also holds: AQI agents might not select plantIED due to focusing on actions that are more prefered across the simulation. These decisions depend on the complex system of agents and their environment, so they cannot be directly known a priori.

4.4.4 Production Example. When the AQI agent decides to perform the plantIED action, the agent restarts this cycle by performing the action where it might be attended by observers. These observers process that event using the processes described in this example, allowing the affordance to spread as a meme through the village.

5 Iraqi Village Simulation Analysis

The Hamariyah Iraqi Village environment models competition between the spread of behaviors: providing intel to the U.S. (giveInformation) and volunteering to help anti-U.S. elements plant an IED on a U.S.-owned building (plantIED). The simulation runs were used to examine three questions about the realism:

1. Diffusion dynamics—Does social learning follow social ties?
2. Cluster formation—Do agents form clusters of adopters?
3. Cluster comparison—What traits determine membership in clusters?

Data from the simulation runs were analyzed to examine each of these questions in Sections 5.1, 5.2, and 5.3 respectively.

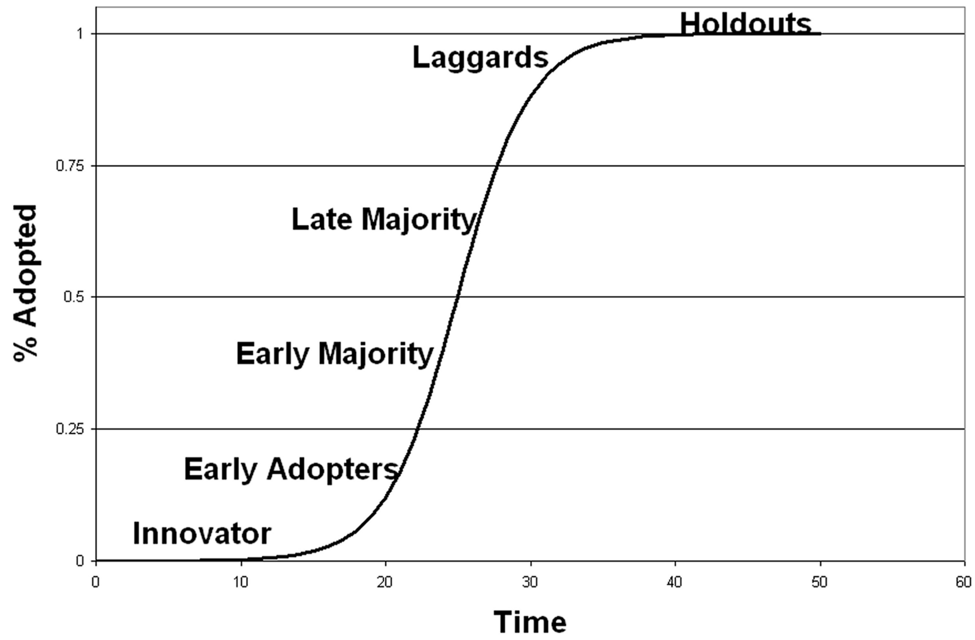

Figure 6. *Diffusion of innovations curve.*

5.1 Diffusion Dynamics

The simulation dynamics give an overview of how the behaviors spread. Behaviors spread in two phases: learning the affordance and expressing the action. Both behaviors spread quickly enough to approach equilibrium within the simulation time horizon, as noted earlier in Section 4.2. The learning curve of each behavior follows a punctuated version of the Rogers (1995) diffusion of innovations process, shown in Figure 6. These patterns indicate a progression of innovators, early adopters, early majority, late majority (late adopters), and laggards. Holdouts are individuals who never adopt and cause the curve to saturate at less than 100% adoption.

Figure 7 shows the percentage of each group that learned giveInformation over time, as the mean of the 20 runs done in this condition. The x-axis shows the

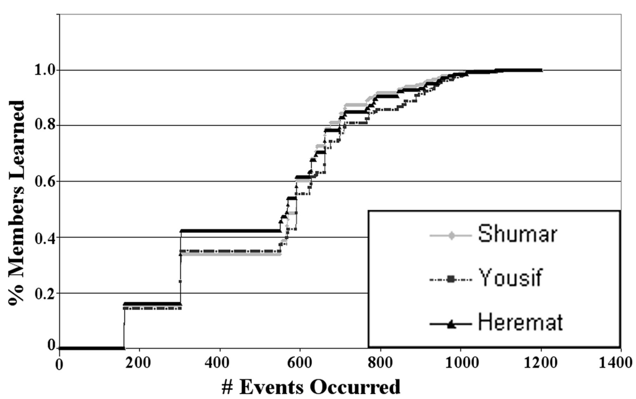

Figure 7. *Percent of group learned giveInformation (hypothesis condition).*

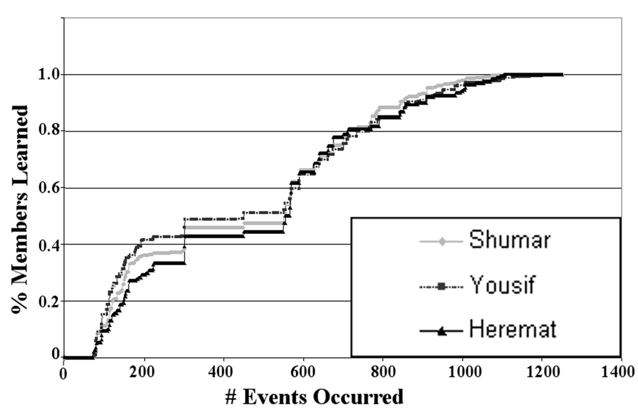

Figure 8. *Percent of group learned giveInformation (randomized condition).*

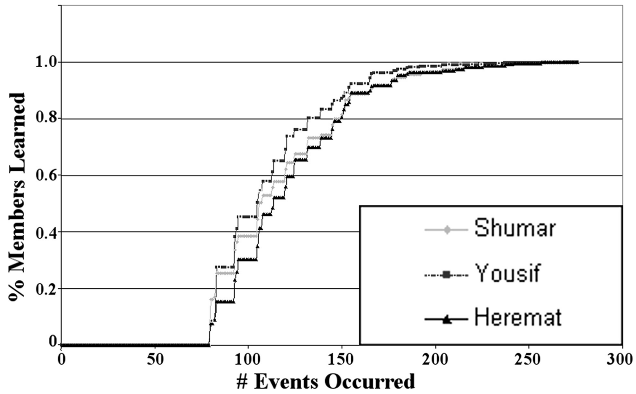

Figure 9. *Percent of group learned plantIED (hypothesis condition).*

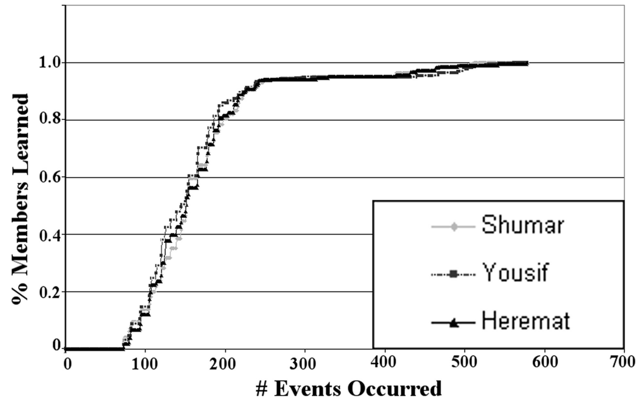

Figure 10. *Percent of group learned plantIED (randomized condition).*

total number of all events that occurred within the simulation, which correlates with time passing. Events are used, because agents can only learn by observing some event. The y-axis shows the fraction of agents who have learned the behavior. To avoid bias from the initial set of agents aware of the behavior, this chart only considers agents who did not start out by knowing the behavior. To help examine the learning region, this chart is truncated at the point where saturation was typically reached (all agents aware of the action). Next to it, Figure 8 shows this same statistic for the randomized condition.

Comparing Figure 7 and Figure 8, it is evident that changing the initial set of agents changes the learning curve of each group. Under the hypothesis condition, giveInformation is initially known by a significant number of Heremat agents. Due to this initial advantage, other Heremat agents tend to learn the behavior faster. In the randomized condition, this learning advantage reverses and the Yousif group members and the Shumar group members have advantages in learning giveInformation. In both conditions, the difference in learning only holds through the early adopter and early majority phases. Once the late majority phase starts, no particular group shows a significant advantage. Despite which group has an advantage, the diffusion rate is fairly similar—reaching saturation after approximately the same number of events.

The same comparison is shown for the plantIED action, shown in Figure 9 (hypothesis) and Figure 10 (randomized). In both conditions, the Yousif group had an advantage in learning rate. For the hypothesis condition, a significant number of the initial carriers are members of the Yousif group. This allows them to better spread the behavior among their own group. In the randomized condition, the Yousif group was also slightly favored in learning plantIED. This indicates that the Yousif are in general more likely to learn this behavior. Additionally, the rate of learning plantIED was greatly impacted by the starting condition. When given to a random set of agents, learning takes twice as long to saturate the population and diffusion is more homogenous across groups (e.g., very similar curves in Figure 10). It is also slower during the steeper part of the learning curve, consistently lagging behind. This means that the starting set for plantIED is more successful in getting awareness of that meme to the population than a random subset of agents would be.

In general, plantIED was also learned faster than giveInformation. This may be due to plantIED occurring more frequently than giveInformation. A t-test was run to test for the probability that there were more plantIED actions than giveInformation actions for both experimental conditions. The t-test strongly indicates plantIED was more common than giveInformation ($p < .01$, 19 DOF). A second t-test also confirmed that the hypothesis condition has a higher frequency of plantIED than the randomized condition, explained by the hypothesis innovators being more likely to perform plantIED than a random set of agents.

Table 5. *Demographic Properties for Cross-Cluster Analysis*

Property	Data type	Description
Group valences		
Valence(U.S. group)	Continuous	Like/dislike toward the U.S. group
Valence(Heremat group)	Continuous	Like/dislike toward the Heremat group
Valence(Shumar group)	Continuous	Like/dislike toward the Yousif group
Valence(Yousif group)	Continuous	Like/dislike toward the Yousif group
Group memberships		
Member of Heremat	Dichotomous	True only if agent in Heremat faction
Member of Shumar	Dichotomous	True only if agent in Shumar faction
Member of Yousif	Dichotomous	True only if agent in Yousif faction
Social properties		
Authority	Continuous	Authority of the agent in his/her group
EmploymentLevel	Dichotomous	If true, agent is employed and typically goes to work during the day
GSP personality factors	Continuous	Personality traits

5.2 Cluster Formation of Adopters

This section examines whether agents formed clusters of adopters, and what characteristics made these clusters distinct. Agents were classified based on two adoption factors: average time of first learning, and average time of first expression. Since agents may not learn or express the behavior, not all agents or clusters have a numeric time value. When an agent was a nonadopter, the learning and/or expression time value was technically classified as never during that run. However, averages and charts require numeric values. Rather than exclude nonadopters from such analyses, time values of never were replaced by the final simulation step (step 3,456). As such, agents and clusters displayed as adopting on the final time step should be considered nonadopters.

Using these factors, the mClust clustering algorithm was used to generate an optimal set of Gaussian expectation–maximization clusters based upon the pair of variables (Fraley & Raftery, 2003). Gaussian clustering was applied for this purpose because the clusters formed fairly regular elliptical patterns that were well-classified using this technique. Clusters will be referred to by their central means during the discussion, in the form of (first learning time, first expression time). Also, it should be noted that some clusters appear as lines when charted because they have little variance on one axis. This is most notable for the holdout and early adopter clusters that form at the edges when graphed.

5.2.1 Cross-Cluster Analysis. To examine the differences between these clusters, a set of demographic properties was collected from the agents belonging to each cluster. The set of properties used for clustering is shown in Table 5. These properties include GSP personality factors, group memberships, valences toward other groups, authority, and employment level. These factors are introduced in Section 3.6.1, Section 3.4.8, Section 3.4.7, Section 3.4.4, and Section 4.1, respectively. Due to the large number of GSP nodes, each node will be briefly described in-text if it has a particular significance for analysis; alternately, the reader can refer back to Table 2 for the full set. Sections 5.2.2 and 5.2.3 examine cluster formation at a high level and focus on individual cluster characteristics.

Clusters were contrasted against other clusters in the same condition. For continuous properties, a one-way ANOVA was run to detect any significant differences

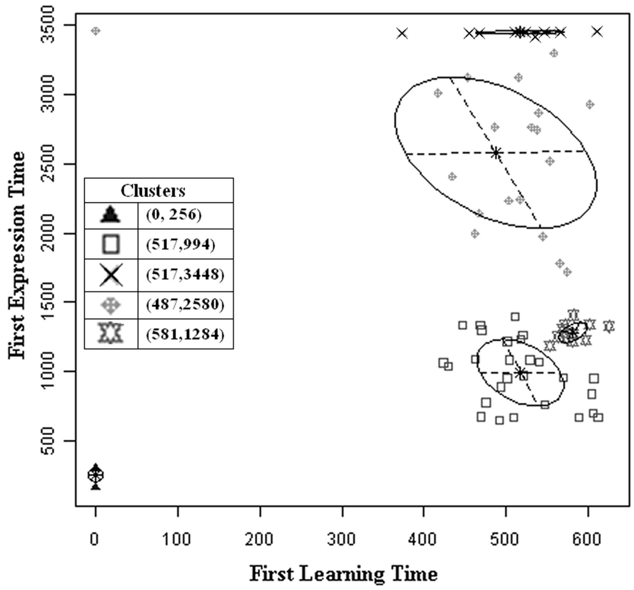

Figure 11. *GiveInformation clusters (hypothesis condition).*

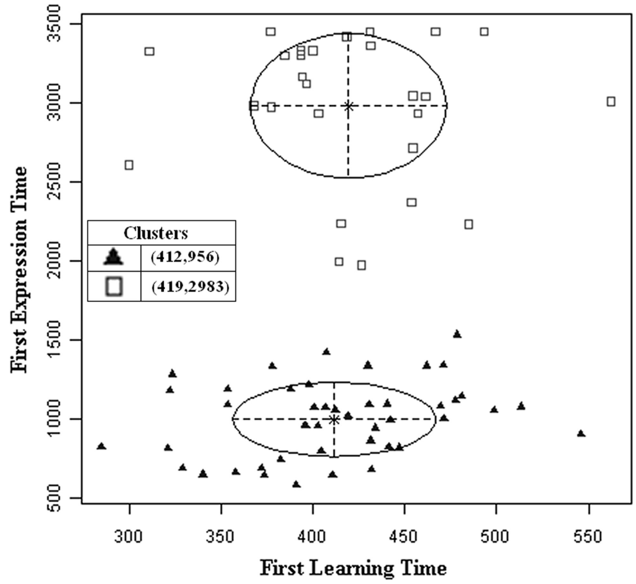

Figure 12. *GiveInformation clusters (randomized condition).*

between clusters. For dichotomous variables, a χ^2 test was run to detect significant differences. After this, a Scheffe post hoc test was applied to examine the specific differences between individual clusters. A very large number of differences were significant, $p < .05$, so only key identifiers that were most unique to each cluster will be discussed. Each key identifier was significant at the .05 level in differentiating a particular cluster, based upon the Scheffe post hoc test. Section 5.3 contains the key indicators for learning and adoption that were discovered through cross-cluster analysis.

5.2.2 GiveInformation Cluster Formation. The clustering results for giveInformation in the hypothesis and randomized conditions are shown in Figure 11 and Figure 12, respectively. The difference between these conditions is different not only in the members of the clusters, but in the number of clusters overall. The hypothesis condition shows five clusters, while the random condition shows only two.

The cluster in the lower left hand (0, 256) is the initial set of agents aware of the behavior, who tend to express it relatively early. At the upper right hand of the graph (517, 3448) is a significant number of agents who learn giveInformation late and most never express it. Of the remaining three clusters, those centered at (517, 993) and (487, 2580) were diffuse but (581, 1284) was very dense. The randomized condition was much simpler—containing only two diffuse groups for learning and expression located at (412, 956) and (419, 2983). Interestingly, both clusters have similar learning times, but very different expression times.

Table 6 shows the basic information about each cluster in the hypothesis condition, including its size and dominant groups represented. Also, each cluster is categorized into its adoption category. One notable category is the holdout set. These agents generally did not express the behavior at all. In this respect, they were not laggards but were typically holdouts for giveInformation. The randomized condition washed out most of these clusters, with later adopters and holdouts in one cluster and early adopters in another cluster.

5.2.3 PlantIED Cluster Formation. PlantIED shared some similarities in its learning and first expression dynamics. Figure 13 and Figure 14 show the mClust cluster graphs for plantIED for the hypothesis and randomized conditions, respectively. As with the giveInformation, the hypothesis condition showed cleaner clusters than the randomized condition.

Table 6. *Demographics for giveInformation Clusters (Hypothesis)*

Cluster location	Cluster size	Primary groups	Learning adoption	Expression adoption
(0, 256)	3	Heremat	Innovator	Early adopter
(517, 994)	27	Shumar	Late majority	Early majority
(581, 1284)	12	Yousif	Laggard	Late majority
(487, 2580)	19	Shumar, Heremat	Early majority	Laggard
(517, 3448)	11	Yousif, Shumar	Late majority	Holdout

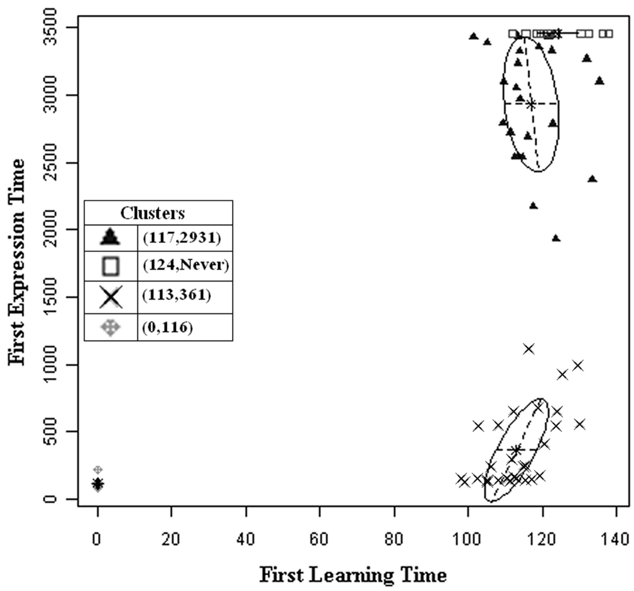

Figure 13. *PlantIED clusters (hypothesis condition).*

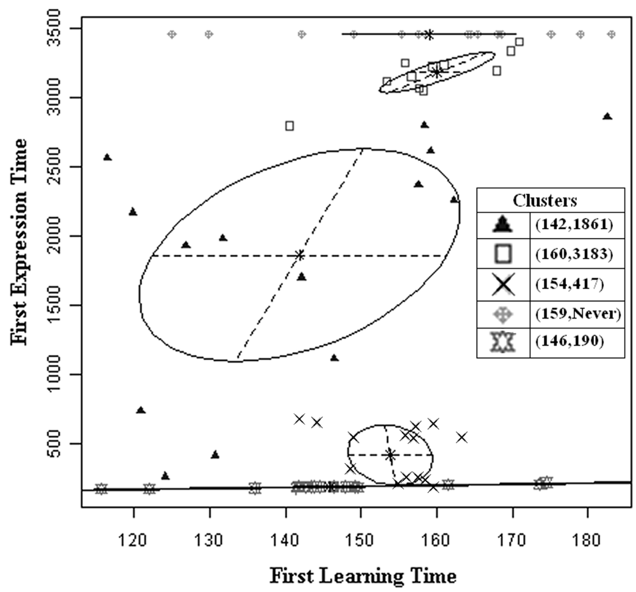

Figure 14. *PlantIED clusters (randomized condition).*

Table 7 shows the basic demographics for the hypothesis clusters and their approximate adoption positions. Even more so than giveInformation in the hypothesis condition, the clusters closely correlate with group membership. The majority of Shumar and Heremat learn the behavior later and wait much longer to express it, if at all. Conversely, a subset of the Shumar and Yousif quickly move to express the behavior. PlantIED is interesting in this condition because learning and first expression track each other quite closely. The agents who are last to learn this behavior are also the least likely to want to express it. This is at a contrast with giveInformation, where expression holdouts still learned it at about the same rate as other agents. In this case, attention correlates well with the motivation to imitate.

The randomized condition for plantIED shows interesting behavior. Table 8 shows the basic demographics for the plantIED action under the randomized condition. While giveInformation was reduced to two clusters, plantIED still displays five clusters in the randomized condition. However, these are not the same five clusters. For many of the cases, this is a small reshuffling, but some agents expressed at different times due to the changes in learning patterns. For example, the hypothesis cluster at (117, 2931) breaks into two smaller clusters. One of those clusters (160, 3183) has a much

Table 7. *Demographics for plantIED Clusters (Hypothesis Condition)*

Cluster location	Cluster size	Primary groups	Learning adoption	Expression adoption
(0, 116)	6	Yousif, Shumar	Innovator	Early adopter
(113, 361)	30	Shumar, Yousif	Early adopter	Early majority
(117, 2931)	23	Shumar, Heremat	Early majority	Late majority
(124, Never)	13	Shumar, Heremat	Late majority	Holdout

Table 8. *Demographics for plantIED Clusters (Randomized Condition)*

Cluster location	Cluster size	Primary groups	Learning adoption	Expression adoption
(146, 190)	17	Yousif, Shumar	Early majority	Early adopter
(154, 417)	15	Shumar, Yousif	Late majority	Early majority
(142, 1861)	14	Shumar, Heremat	Early majority	Late majority
(160, 3183)	11	Mixed	Late majority	Laggard
(159, Never)	15	Shumar, Heremat	Late majority	Holdout

longer amount of time before first expression, while the other includes four additional subjects (142, 1861).

While randomized innovators compress the differences in learning, small differences persist for plantIED. The randomized condition shows a correlation between the time that a cluster learns and when it adopts the action. This means that agents who are more likely to perform plantIED also learn it quicker, regardless of who initially spreads the behavior. These differences are likely due to factors such as motivated attention and differences in group size.

5.3 Cross-Cluster Comparison Results

The prior analysis demonstrated that distinct clusters of agents exhibited different learning and adoption patterns. This section examines the identifying features that differentiated particular clusters, based on the ANOVA and Scheffe tests described in Section 5.2.1.

5.3.1 GiveInformation Cluster Identifying Features. For giveInformation, group membership and GSP factors were the strongest determinants of cluster membership. In the hypothesis condition, the clusters can be thought of as following three main behavioral patterns: innovators (0, 256), holdouts (517, 3448), and fence-sitters (the middle three clusters). The innovator cluster at (0, 256) was small and not very influential. Notably, it does not include all the agents who start with the giveInformation action. This cluster has high valence toward the U.S. group, while all other clusters have low valence toward the U.S. The innovators are mainly Heremat, but most Heremat members are part of the (487, 2580) cluster, making them fence-sitters and late adopters. These agents are some of the first ones to learn the behavior, but among the last to try it. One of the differences between the innovator group and this cluster is that the innovators give a higher weight to be_relationship_focused in their GSP.

The holdout cluster at (517, 3448) lies at the opposite end of the spectrum. Intuitively, one might assume that the holdouts dislike the U.S. group. Intuition would be wrong: the holdouts are not very different from the fence-sitting clusters in their group membership or valence. The ANOVA analysis indicates that holdouts place a very high value on personal interests and safety (i.e., GSP weights for high safety and for_the_self). They also have a much lower inclination to control their environment, as shown by low importances for esteem and be_controlling nodes. Overall, this cluster of agents shares a personality type that is not inclined to take risks. Considering that becoming an informant is a dangerous endeavor, these agents would simply rather stay home.

The fence-sitting clusters for the hypothesis condition differ mainly by group membership. (517, 994) is a Shumar-dominated group, (581, 1284) is a Yousif-dominated group, and (487, 2580) is mostly Heremat. The difference in learning times is explained by the relationship of each cluster and the Heremat group, who dominate the innovators. These clusters largely disappear in the randomized condition, where initial knowledge is randomized. The meaningful social patterns seen in the hypothesis condition disappear in this condition. Without the initial social biases, behavior is learned across groups more evenly.

5.3.2 PlantIED Cluster Identifying Features. PlantIED had a sharp distinction between each cluster in the hypothesis condition. These four groups can be thought of as innovators, would-be innovators, late majority, and holdouts. The (0, 116) innovators of the plantIED action were prone to expressing the action because they felt it would benefit their group's future (for_the_group, own_group), as well as to satisfy their esteem goals and assert_individuality standards. They also place a low weight on safety. They are also primarily Yousif group members, and share a negative valence toward the U.S. group.

The early adopters at (113, 361) can be considered to be would-be innovators, due to their strong similarity with the agents in (0, 116). They mainly differ because they need to learn the affordance before performing it. Their long-term preferences are oriented toward symbolic nodes rather than materialism or for_own_group nodes. This difference appears to be the influence of AQI members in the cluster. Overall, these agents waste little time between learning the affordance before volunteering to plant an IED.

The late majority and holdouts are distinct. The (117, 2931) cluster, which is partially resistant to expressing the action, is business-oriented and pro-U.S. It places high importance on growing economic resources (materialism), conforming to society (conform_to_society), and positive outcomes for the self (for_the_self). It also places a higher importance on safety than the IED-active clusters, but not as high as the other resistant cluster at (124, Never). It is also the only cluster that conclusively has a high valence toward the U.S. group.

The holdouts at (124, never) are self-interested good guys. They have many good-guy personality traits and are less materialistic and have a high value on for_the_greater_good. However, key primary identifying characteristics are high value to safety and respect_for_life. As a result, this cluster has a major overlap with the resistant agents for giveInformation. Overall, the holdouts have low valence toward the U.S. but are simply unwilling to take risky actions.

The randomized condition shifts the identifying features of the cluster slightly. For example, the cluster at (142, 1861) has a higher employment level and authority level compared to other groups. Additional work responsibilities may play a role in that subgroup's delay in first expression. Most of the prior indicators of early or late first expression still hold. The next section summarizes the indicators that were reliable for both the hypothesis condition and the randomized condition, for learning and first expression times. These will be referred to as the key indicators for the type of agent and the situation that leads agents to adopt giveInformation or plantIED.

5.3.3 Key Indicators: Summary. A summary of the key indicators that differentiated early learners versus late learners is listed in Table 9, for giveInformation and plantIED. From this analysis, the early learners were differentiated primarily by their social cues (e.g., in-group, reference groups, valence) which account for most of the variance in learning. Motivated attention was a secondary influence on top of this for plantIED. Attention cues such as novelty, selective attention, and transferability did not strongly influence learning times between clusters. Primarily, these factors were not indicators because they can vary over an individual agent's trajectory rather than differing greatly between agents. As such, patterns of social learning in this virtual

Table 9. *Key Indicators for Determining Social Learning*

Key indicator	GiveInformation learning time change	PlantIED learning time change
Same ingroup as innovators	Faster learning	Faster learning
High valence toward innovator's group	Slightly faster learning	Slightly faster learning
Low motivated attention to action	No clear connection	Slightly slower learning
Innovators express earlier	Faster learning	Faster learning

environment reflect the preexisting structure of social cues.

The randomized condition cases showed that a purely random subset of innovators significantly equalizes the learning rates, on average. This confirms that social cues (who does the behavior) provide more consistent indicators than central cues (what kind of behavior occurred). Finally, all agents learned earlier if the innovators expressed earlier. Since agents cannot learn about the new behavior except when other agents express it, this relationship was expected.

On the converse, early expression of behavior is dominated by personality factors. Table 10 shows the key indicators that help determine whether an agent will express a behavior earlier or later. Valence toward the U.S. group is the only consistent nonpersonality key factor that influences expression of either behavior in this simulation. Employment may have also been an environmental influence that delayed plantIED, but was not consistently statistically significant. Otherwise, expression was almost entirely determined by the personality factors. Safety goals were a key limiting factor for both behaviors, an obvious connection for dangerous actions. However, seemingly unrelated factors such as long-term preferences for_the_self and materialism had a significant influence as well. This indicates that these behaviors are competing with day-to-day activities and pursuing economic endeavors.

6 Discussion

Experiments with Hamariyah demonstrated the feasibility of representing realistic adoption patterns of new behavior in a virtual world. Rather than losing control over the virtual environment, agents produced well-formed patterns of learning and adoption of behavior. These patterns were produced by a double-compete process that mediated the spread of behavior. Competition at the attention level produced learning patterns that were based on social cues and motivated attention to the behavior. Competition at the decision level produced patterns in adopting each new behavior based on an agent's social ties (group valence) and motivation (personality factors).

6.1 Realistic Spread of Behavior

The NonKin simulation fundamentally works as a complex system, with significant probabilistic and path-dependent effects on adoption. Nominally, these effects prevent predictive and repetitive behavior that users cite as a problem with virtual agents (Bickmore et al., 2010). However, users of a virtual environment will lose their sense of immersion if the overall patterns of adoption fail to follow reasonable patterns. Three aspects of the patterns will be discussed here: diffusion, cluster detection, and cluster prediction.

The diffusion patterns indicated that the behaviors spread plausibly through the population. The high-level dynamics demonstrated a punctuated version of the adoption curve expected for diffusion of innovations (Rogers, 1995). These patterns also indicate that the spread was more realistic when the initial innovators were chosen based on their personal traits and social ties, rather than chosen randomly. This indicates that social learning effectively helped agents transmit behavior within their group and to friendly agents. These patterns

Table 10. *Key Indicators for First Expression (Adopting Behavior)*

Key indicator	GiveInformation first expression time	PlantIED first expression time
↑ Valence toward U.S.	Earlier expression	Later expression (or none)
GSP goals (short-term values)		
↑ Safety	Prevents expression	Prevents expression
↑ Esteem	—	Earlier expression
GSP standards (preferred methods)		
↑ Assert individuality	Earlier expression	Later expression
↑ Be task focused	Later expression	Earlier expression
↑ Be relationship focused	Earlier expression	Later expression
↑ Be controlling	Earlier expression	Earlier expression
↑ Bring about greater good	Earlier expression	—
↑ Use asymmetric attacks	—	Earlier expression
GSP preferences (long-term wants)		
↑ For own group	—	Earlier expression
↑ For the self	Later expression	Later expression
↑ Materialistic	Later expression	Later expression
↑ Symbolic	Earlier expression	Earlier expression

also indicate that plantIED was a more popular action than giveInformation. As noted previously in Figure 5, most agents in the simulation have a low valence toward U.S. group so negative actions against them should be more common. Agents were also slightly more likely to learn actions they preferred, showing the influence of motivated attention.

Analysis to detect adoption clusters also yielded a positive result: agents produced well-formed clusters of behavior based on their social influences, personality, and context (e.g., employment level). These clusters represent clear patterns of early versus late adoption, as well as innovators and holdouts. Agents display reasonable patterns of agents gravitating toward either giveInformation or plantIED, as well as patterns of holdouts avoiding both behaviors entirely. Comparison of the hypothesis and randomized conditions demonstrated that meaningful selection of the behavior innovators makes these clusters better defined and more plausible.

Finally, statistically comparing the clusters produced indicators with predictive value about an agent's cluster. These indicators are based on the scenario's initial values rather than its runtime values, which may change over time. The key indicators for these patterns had a high degree of face-validity, such as holdouts being unwilling to take dangerous actions. Identifying key indicators means that a virtual society designer can predict when different agents would learn and adopt a behavior. For a larger multi-agent environment, a smaller set of representative agents can be evaluated to examine the cluster indicators. These indicators can be used to classify new agents added to the virtual society, in order to estimate their expected adoption behaviors.

As such, the spread of behavior can be modeled with high fidelity with respect to who learns and adopts new behavior. This analysis showed that the model was effective for representing competition of behaviors spreading within the fictional Hamariyah Iraqi village. It was possible to determine not only the diffusion of each behavior within the population, but also the key identifying factors that determined why agents adopted a given behavior.

6.2 Modeling and Simulation Findings

While these simulations used a fictional village, the cognitive model was also designed to simulate real-world scenarios (Nye, 2012). As such, social simulations based on this model could offer insight into the real-life adoption patterns. Even for the Hamariyah village, the key indicators gave some interesting insights that connect with theories of counterinsurgency.

GiveInformation adoption was associated with a low weight for personal safety and a high weight for relationship-oriented problem solving. However, based on the personality traits used to create the scenario, agents who valued relationship-oriented problem solving also valued their personal safety highly. This implies that adequate security is pivotal to securing informants. This finding is supported by some counterinsurgency analysts, who view security as essential even in a hearts-and-minds campaign (Krepinevich, 2005). Secondly, employment level was not found to be a significant factor for volunteering to participate in IED activities. While work-related tasks might delay volunteering slightly, if agents are willing to risk their lives, then they are also willing to find time to do so. This is concordant with research such as Berman, Felter, and Shapiro (2009), who state that higher employment does not appear to decrease the likelihood of violent rebellion activities that result in civilian deaths.

While it is important not to extrapolate too much from the results of a virtual training scenario, these findings indicate some potential for significant analytical value by applying this approach to real-life scenarios. For this potential to be realized, a village would need to be calibrated and validated using data based on a specific real-life scenario. Additionally, second-order effects such as external influences and communication mediums would be important for studying a village situated in a larger social system. While gathering data for modeling a specific case study is challenging, other projects based on the PMFServ architecture have previously used a system of structured subject-matter experts and databases to generate scenarios for forecasting purposes (Silverman, Bharathy, & Kim, 2009; Bharathy & Silverman, 2010; O'Brien, 2010).

7 Conclusions and Future Directions

This research is part of a larger class of topics that increase realism by focusing on the realistic patterns of a virtual agent society, rather than on an individual agent. This paradigm shift from virtual agents to virtual agent societies is a significant trend within virtual environments. Representing and studying the spread of behavior among virtual agents is an important direction for the realism of immersive environments. Using cognitively-based agents, this work demonstrated that plausible patterns of learning and adoption of behavior can be added to an immersive training environment.

An open question is how to extend this work to model the abandonment of behavior, which has recently been explored by social psychologists (Berger & Heath, 2007). Real societies are dynamic, with new trends and cliques of behavior forming and disbanding. Cognitive models that emulate human abandonment of behavior would be a logical next step for supporting dynamic trends in behavior by virtual agents. Particularly for long-running immersive environments, such as massively multiplayer online (MMO) systems and virtual worlds, extinction of behavior may be of equal importantce as adoption.

This model should also have value for social simulation. Observational learning, multi-layered social cues, and contextual social learning have not been well-examined using social simulations. A significant challenge to such research is the amount of data necessary to initialize detailed agents, who require numerous measures of personality and social relationships. However, given the potential benefits of using social simulation to predict classes of adopters down to the individual level, this direction fills a role not fully addressed by existing approaches.

A final direction is to survey user reactions to these adoption patterns, studying user perceptions of immersion in the NonKin environment. This work demonstrates that cognitive agents can plausibly model adoption patterns, rather than relying on static action sets or simple random patterns. However, the level to which these patterns improve realism and reduce perceived repetition must still be explored. Presence questionnaires and other measures can provide valuable insight into these issues (Witmer & Singer, 1998; Jennett et al., 2008). Perceptions of system efficacy, such as immersion, have been shown to influence performance outcomes in training environments (Jia, Bhatti, & Nahavandi, 2012). Quantifying the impact of behavioral trends on immersion would help define their role in training and gaming environments.

Acknowledgments

Many thanks to the other researchers and staff at the ACASA lab who worked on the NonKin framework, without whom this work would not have been possible. Benjamin Nye would like to thank Barry Silverman, whose vision created the NonKin framework that this work builds upon. Finally, thanks to the U.S. Marine Corps for providing the human terrain data used to develop the Hamariyah village scenario.

References

Arbib, M. A. (2011). From mirror neurons to complex imitation in the evolution of language and tool use. *Annual Review of Anthropology*, *40*, 257–273.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American, 193*(5), 31–35.

Axelrod, R. (1997). Advancing the art of simulation in the social sciences. *Complexity, 3*(2), 16–22.

Bandura, A. (1986). *Social foundations of thought and action: A social cognitive theory*. Englewood Cliffs, NJ: Prentice Hall.

Barrios-Aranibar, D., Alsina, P., Nedjah, N., Coelho, L., & Mourelle, L. (2007). Imitation learning: An application in a micro robot soccer game. In N. Nedjah, L. d. S. Coelho, & L. d. M. Mourelle (Eds.), *Mobile robots: The evolutionary approach. Studies in computational intelligence*, Vol. 50 (pp. 201–219). Berlin: Springer.

Berger, J., & Heath, C. (2007). Where consumers diverge from others: Identity signaling and product domains. *Journal of Consumer Research*, *34*(2), 121–134.

Berman, E., Callern, M., Felter, J., & Shapiro, J. N. (2009). *Do working men rebel? Insurgency and unemployment in Iraq and the Philippines.* NBER Working Paper No. 15547. Cambridge, MA: National Bureau of Economic Research.

Bharathy, G. K., & Silverman, B. G. (2010). Validating agent based social systems models. *Proceedings of the Winter Simulation Conference (WSC) 2010*, 441–453.

Bickmore, T., Schulman, D., & Yin, L. (2010). Maintaining engagement in long-term interventions with relational agents. *Applied Artificial Intelligence*, *24*(6), 648–666.

Billard, A., & Dautenhahn, K. (1999). Experiments in learning by imitation — Grounding and use of communication in robotic agents. *Adaptive Behavior, 7*(3–4), 415–438.

Blair, K., Schwartz, D., Biswas, G., & Leelawong, K. (2007). Pedagogical agents for learning by teaching: Teachable agents. *Educational Technology and Society, 47*(1), 56–61.

Bornstein, R. F. (1989). Exposure and affect: Overview and meta-analysis of research, 1968–1987. *Psychological Bulletin*, *106*(2), 265–289.

Centola, D. (2010). The spread of behavior in an online social network experiment. *Science*, *329*(5996), 1194–1197.

Cherry, C. E. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the Acoustical Society of America, 25*, 975–979.

Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Penguin Books.

Delre, S. A., Jager, W., Bijmolt, T. H. A., & Janssen, M. A. (2010). Will it spread or not? The effects of social influences and network topology on innovation diffusion. *Journal of Product Innovation Management*, *27*(2), 267–282.

Dennett, D. C. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. New York: Simon and Schuster.

Dignum, F. (2012). Agents for games and simulations. *Autonomous Agents and Multi-Agent Systems*, *24*(2), 217–220.

Dignum, V. (2009). *Handbook of research on multi-agent systems: Semantics and dynamics of organizational models.* Hershey, PA: Idea Group Inc. (IGI).

Fazio, R. H., Roskos-Ewoldsen, D. R., & Powell, M. C. (1994). Attitudes, perception, and attention. In P. M. Niedenthal & S. Kitayama (Eds.), *The heart's eye: Emotional influences in perception and attention* (pp. 197–216). New York: Academic Press.

Fraley, C., & Raftery, A. E. (2003). Enhanced model-based clustering, density estimation, and discriminant analysis software: MCLUST. *Journal of Classification*, *20*(2), 263–286.

Gaver, W. W. (1991). Technology affordances. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Reaching Through Technology* (pp. 79–84).

Gibson, J. J. (1986). *The ecological approach to visual perception*. Mahwah, NJ: Lawrence Erlbaum Associates.

Gruhl, D., Guha, R., Liben-Nowell, D., & Tomkins, A. (2004). Information diffusion through blogspace. *Proceedings of the 13th Conference on the World Wide Web*, 491.

Hermann, M. G. (2005). Assessing leadership style: A trait analysis. In J. M. Post (Ed.), *The Psychological Assessment of Political Leaders* (pp. 178–214). AnnArbor, MI: University of Michigan Press.

Hilmert, C. J., Kulik, J. A., & Christenfeld, N. J. S. (2006). Positive and negative opinion modeling: The influence of another's similarity and dissimilarity. *Journal of Personality and Social Psychology*, *90*(3), 440–452.

Hofstede, G. (2003). *Culture's consequences: Comparing values, behaviors, institutions, and organizations across nations.* Thousand Oaks, CA: Sage Publications.

Holland, J. (1998). *Emergence: From chaos to order* (Vol. 2). New York: Perseus Books.

House, R. J., Hanges, P. J., Javidan, M., Dorfman, P. W., & Gupta, V. (2004). *Culture, leadership, and organizations: The GLOBE study of 62 societies.* Thousand Oaks, CA: Sage Publications.

James, W. (1890). *The principles of psychology.* Cambridge, MA: Harvard University Press.

Jennett, C., Cox, A. L., Cairns, P., Dhoparee, S., Epps, A., Tijs, T., & Walton, A. (2008). Measuring and defining the experience of immersion in games. *International Journal of Human–Computer Studies*, *66*(9), 641–661.

Jia, D., Bhatti, A., & Nahavandi, S. (2012). The impact of self-efficacy and perceived system efficacy on effectiveness of virtual training systems. *Behaviour and Information Technology.* Advance online publication. doi: 10.1080/0144929X.2012.681067

Johnston, W. A., Hawley, K. J., Plewe, S. H., Elliott, J. M. G., & DeWitt, M. J. (1990). Attention capture by novel stimuli. *Journal of Experimental Psychology: General*, *119*(4), 397–411.

Kameda, T., Ohtsubo, Y., & Takezawa, M. (1997). Centrality in sociocognitive networks and social influence: An illustration in a group decision-making context. *Journal of Personality and Social Psychology*, *73*(2), 296–309.

Kelley, G. A. (1955). *The psychology of personal constructs.* New York: W. W. Norton.

Kerr, W., Cohen, P. R., & Adams, N. (2011). Recognizing players' activities and hidden state. *Proceedings of Foundations of Digital Games 2011*, 84–90.

Kerr, W., Hoversten, S., Hewlett, D., Cohen, P., & Chang, Y.-H. (2007). Learning in Wubble World. *IEEE Conference on Development and Learning*, 330–335.

Knox, W. B., Fasel, I., & Stone, P. (2009). Design principles for creating human-shapable agents. *Proceedings of the AAAI Spring 2009 Symposium on Agents that Learn from Human Teachers.*

Krepinevich A. F., Jr. (2005). How to win in Iraq. *Foreign Affairs*, *84*(5), 87–104.

Laird, J. E. (2008). Extending the Soar cognitive architecture. *Proceedings of the 2008 Conference on Artificial General Intelligence*, 224–235.

Lee, D. K., Itti, L., Koch, C., & Braun, J. (1999). Attention activates winner-take-all competition among visual filters. *Nature Neuroscience*, *2*(4), 375–381.

Li, X., Mao, W., Zeng, D., Wang, F.-Y., Yang, C., Chen, H., … Zhan, J. (2008). Agent-based social simulation and modeling in social computing. In C. C. Yang et al. (Eds.), *Intelligence and security informatics. Lecture notes in computer science*, Vol. 5075 (pp. 401–412). Berlin: Springer.

Lim, H. C., Stocker, R., Barlow, M., & Larkin, H. (2011). Interplay of ethical trust and social moral norms: Environment modelling and computational mechanisms in agent-based social simulation (ABSS). *Web Intelligence and Agent Systems*, *9*(4), 377–391.

Mackintosh, N. J. (1983). *Conditioning and associative learning.* New York: Oxford University Press.

Mantell, D. M. (1971). The potential for violence in Germany. *Journal of Social Issues*, *27*(4), 101–112.

Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, *50*, 370–396.

Milgram, S. (1963). Behavioral study of obedience. *The Journal of Abnormal and Social Psychology*, *67*(4), 371–378.

Nye, B. D. (2011). *Modeling memes: A memetic view of affordance learning.* Ph. D. thesis, University of Pennsylvania, Philadelphia, PA.

Nye, B. D. (2012). Modeling socially transmitted affordances: A computational model of behavioral adoption tested against archival data from the Stanford Prison Experiment. *Proceedings of the Behavior Representation in Modeling and Simulation Conference (BRIMS).* Orlando, FL: SISO.

O'Brien, S. P. (2010). Crisis early warning and decision support: Contemporary approaches and thoughts on future research. *International Studies Review*, *12*(1), 87–104.

Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions.* Cambridge, UK: Cambridge University Press.

Panait, L., & Luke, S. (2005). Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, *11*(3), 387–434.

Pareto, L., Arvemo, T., Dahl, Y., Haake, M., Gulz, A., Biswas, G., … Mitrovix, A. (2011). A teachable-agent arithmetic game's effects on mathematics understanding, attitude and self-efficacy. In G. Biswas, S. Bull, J. Kay, & A. Mitrovic (Eds.), *Artificial intelligence in education. Lecture notes in computer science*, Vol. 6738 (pp. 247–255). Berlin: Springer.

Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Advances in Experimental Social Psychology, 19*, 123–205.

Platow, M. J., Haslamb, S. A., Botha, A., Chewa, I., Cuddona, M., Goharpeya, N., … Grace, D. M. (2005). "It's not funny if they're laughing": Self-categorization, social influence, and responses to canned laughter. *Journal of Experimental Social Psychology*, *41*(5), 542–550.

Railsback, S. F. (2001). Concepts from complex adaptive systems as a framework for individual-based modelling. *Ecological Modelling*, *139*(1), 47–62.

Rogers, E. M. (1995). *Diffusion of innovations.* New York: Free Press.

Rogers, E. M., Medina, U. E., Rivera, M. A., & Wiley, C. J. (2005). Complex adaptive systems and the diffusion of innovations. *The Innovation Journal: The Public Sector Innovation Journal, 10*(3), Article 30.

Schreiber, C., & Carley, K. M. (2007). Agent interactions in Construct: An empirical validation using calibrated grounding. *Proceedings of the Behavior Representation in Modeling and Simulation Conference (BRIMS)*. Orlando, FL: SISO.

Shannon, C. E. (1948). A mathematical theory of communication. *Key Papers in the Development of Information Theory.* Retrieved from cm.bell-labs.com /cm/ms/what/shannonday/shannon1948.pdf

Silverman, B. G., & Bharathy, G. K. (2005). Modeling the personality and cognition of leaders. *Proceedings of the Behavior Representation in Modeling and Simulation Conference (BRIMS)*. Orlando, FL: SISO.

Silverman, B. G., Bharathy, G. K., Johns, M., Eidelson, R. J., Smith, T. E., & Nye, B. D. (2007). Socio-cultural games for training and analysis. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans*, *37*(6), 1113–1130.

Silverman, B. G., Bharathy, G. K., & Kim, G. J. (2009). Challenges of country modeling with databases, newsfeeds, and expert surveys. In A. Uhrmacher & D. Weyns (Eds.), *Multi-agent systems: Simulation and applications* (pp. 271–300). New York: Taylor and Francis.

Silverman, B. G., Bharathy, G. K., Nye, B. D., & Eidelson, R. J. (2007). Modeling factions for "Effects based operations": Part I, Leader and follower behaviors. *Journal of Computational and Mathematical Organization Theory*, *13*(4), 379–406.

Silverman, B. G., Johns, M., Cornwell, J. B., & O'Brien, K. (2006). Human behavior models for agents in simulators and games: Part I: Enabling science with PMFserv. *Presence: Teleoperators and Virtual Environments*, *15*(2), 139–162.

Silverman, B. G., Pietrocola, D., Nye, B. D., Weyer, N., Osin, O., Johnson, D., & Weaver, R. (2012). Rich socio-cognitive agents for immersive training environments—Case of NonKin Village. *Autonomous Agents and Multi-Agent Systems*, *24*(2), 312–343.

Silverman, B. G., Pietrocola, D., Weyer, N., Weaver, R., Esomar, N., Might, R., & Chandrasekaran, D. (2009). NonKin village: An embeddable training game generator for learning cultural terrain and sustainable counter-insurgent operations. *Agents for Games and Simulations, Lecture notes in artificial intelligence,* Vol. 5920 (pp. 135–154). Berlin: Springer.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28*(9), 1059–1074.

Sun, R. (2007). Cognitive social simulation incorporating cognitive architectures. *IEEE Intelligent Systems*, *22*(5), 33–39.

Tajfel, H. (1982). Social psychology of intergroup relations. *Annual Reviews in Psychology*, *33*(1), 1–39.

Tanford, S., & Penrod, S. (1984). Social influence model: A formal integration of research on majority and minority influence processes. *Psychological Bulletin*, *95*(2), 189–225.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.

van Eck, P. S., Jager, W., & Leeflang, P. S. H. (2011). Opinion leaders' role in innovation diffusion: A simulation study. *Journal of Product Innovation Management, 28*(2), 187–203.

Van Segbroeck, S., Jong, S. de, Nowe, A., Santos, F. C., & Lenaerts, T. (2010). Learning to coordinate in complex networks. *Adaptive Behavior*, *18*(5), 416–427.

Witmer, B. G., & Singer, M. J. (1998). Measuring presence in virtual environments: A presence questionnaire. *Presence: Teleoperators and Virtual Environments*, *7*(3), 225–240.

Zentall, T. R. (2007). Imitation: Definitions, evidence, and mechanisms. *Animal Cognition*, *9*(4), 335–353.