

Preparing for the Aftermath: Using Emotional Agents in Game-Based Training for Disaster Response

Donna D. Djordjevich, Patrick G. Xavier, Michael L. Bernard, Jonathan H. Whetzel,
Matthew R. Glickman, and Stephen J. Verzi

Abstract: *Ground Truth*, a training game developed by Sandia National Laboratories in partnership with the University of Southern California GamePipe Lab, puts a player in the role of an Incident Commander working with teammate agents to respond to urban threats. These agents simulate certain emotions that a responder may feel during this high-stress situation. We construct psychology-plausible models compliant with the Sandia Human Embodiment and Representation Cognitive Architecture (SHERCA) that are run on the Sandia Cognitive Runtime Engine with Active Memory (SCREAM) software. SCREAM's computational representations for modeling human decision-making combine aspects of ANNs and fuzzy logic networks. This paper gives an overview of *Ground Truth* and discusses the adaptation of the SHERCA and SCREAM into the game. We include a semiformal description of SCREAM.

I. INTRODUCTION

Emerging modes of attack using weapons of mass destruction (WMD), now defined within threat scenarios from the United States Department of Homeland Security (DHS) and the Department of Defense (DoD), require new approaches to examining detection, mitigation, and response options. In recent years, DHS and other government agencies, charged with preparing for WMD attacks and other catastrophic events, have turned to large multi-person exercises using computer-based simulation to address preparedness training. While this approach is extremely valuable, it also suffers from drawbacks: a large number of trainees must use the system at the same time and each threat scenario takes a day or more to complete.

To mitigate those drawbacks, we investigated different methodologies that would lead us to a complimentary solution. Our selected option was to develop a software gaming platform specifically designed to prepare decision makers for WMD attacks. Two key features of this platform are: (a) allowing for single-player training and (b) shorter scenario run-times. We selected 20 minute run-times as our target time limit for the training scenarios to focus our design by abstracting out the details to keep the scenario focused on key learning objectives. This allows us to more effectively

deliver quality training to the target audience at a frequency that facilitates learning the material in a timely manner. Furthermore, the single player mode lets a responder train anytime, anywhere with software agents taking the place of the responder's human counterparts.

However, the nature of the scenarios we are targeting prevents us from simply replacing the response team members with traditional drone-like Non-Player Characters (NPCs). This is because the training scenarios are based on events that trigger emotionally-biased actions from the many people involved. For example, an Incident Commander (IC) in charge of responding to a toxic industrial chemical spill must protect the lives of the general population while keeping their first responders out of harm's way. But what if a trainee orders a police unit to shelter-in-place a city block that is covered in the toxic fog? Traditional drone-like NPCs would blindly obey the command given by the IC trainee and conduct the action regardless of the effects on its safety. Instead, we wish to provide trainees with virtual teammates that might or might not obey the command due to their emotions and the IC trainee's reactions to the incident. In our case, the NPC would detect the hazardous fog and resist the action. Trainees who then force the units to perform the action, regardless of the feedback, are then penalized in multiple ways: the unit suffers from the physical damage, the trainee's response effectiveness score is reduced, and emotional distress is caused to the unit which reduces its ability or interest to perform other requested actions. As such, including emotion into the virtual teammates provides a more accurate representation on the effect of a trainee's decisions in these types of situations.

This paper covers our work in developing emotion models into the virtual teammates for *Ground Truth: Toxic City*, a game-based platform for training Incident Commanders on response strategies following a toxic chemical spill. The paper will detail the Sandia Human Embodiment and Representation Cognitive Architecture (SHERCA), the design used for constructing how fundamental emotions influence the human cognitive process, and how SHERCA is integrated into the virtual teammates built for the training environment.

II. GROUND TRUTH: TOXIC CITY

A. Gameplay Description

Ground Truth: Toxic City is developed leveraging the open-source graphics rendering engine OGRE. We required full source access to include support for our research and

Donna D. Djordjevich is with Sandia National Laboratories, Livermore, CA 94551-9406 (e-mail: dddjord@sandia.gov).

Patrick G. Xavier is with Sandia National Laboratories, Albuquerque, NM 87185-1004 (e-mail: pgxavie@sandia.gov).

Michael L. Bernard is with Sandia National Laboratories, Albuquerque, NM 87185-1188 (email: mlberna@sandia.gov).

Jonathan H. Whetzel, Matthew R. Glickman, and Stephen J. Verzi are with Sandia National Laboratories, Albuquerque, NM 87185-1011 (e-mail: {jhwhezt, mrglick, sjverzi}@sandia.gov).

development objectives. Our game is modeled after the Real-Time Strategy class of gameplay. The trainee assumes the role of the IC and has an overhead view of the city s/he is chartered to protect. A common feature in these types of games is the “fog of war” that represents what the trainee can or cannot see. In our case, we allow the player to see through this fog based on his/her unit’s field of view as shown in Figure 1. As units travel the space, the map shows the ground truth information while the information on darkened space is considered unknown.



Figure 1: Police unit Fog of War

The trainee’s main objective is to save as many lives as possible within the game’s time limit by keeping them away from the toxic cloud. Civilians can only survive for a few minutes inside the cloud area before they are considered significantly injured and treated as a “loss” for sake of scoring. For directing the civilians to safety, the trainee controls a selection of fire, police, and HazMat units. The trainee can: 1) evacuate city blocks, which moves people from buildings onto the city streets, 2) shelter-in-place city blocks, which provides reduced exposure levels, 3) barricade and direct traffic, which affects the throughput of the city streets, and 4) attempt to contain the leak. The player can monitor street traffic on the display with shaded arrows indicating traffic direction and colored regions representing traffic density. Blue colored regions equal low, yellow equals moderate, and red equals severe traffic density. Each unit also has different actions they can perform based on their type. For example, firefighters can don Personal Protective Equipment (PPE) to protect themselves when entering the toxic cloud area.

B. Game AI Virtual Teammates

In *Ground Truth*, agents serve as the virtual teammates the player must direct in order to win the game. Initially, we implemented a reactive agent design for the virtual teammates. The agents executed commands ordered by the player and report on events of interest to the player in game (e.g., location of the fog cloud). These agents did not have any belief-desire-intent (BDI) modeling [1] that would result

in them acting autonomously. This decision on the agent design resulted from the original intent of the game, training first responder commanders on the best methods on where to direct resources to minimize civilian casualties. However, the reactive agent behavior creates a game that oversimplifies the scenario from a realistic training perspective and feedback from subject matter experts made us reconsider this approach.



Figure 2: Ground Truth gameplay view

We developed an agent architecture for the game that would support not only this reactive agent design, but would also allow for the creation of more complex agents that present the player with challenges more in line with the actual environment. The architecture decomposes each of the game actions into states. For *Ground Truth*, the states include MoveTo, Evacuate, Shelter-In-Place, Barricade, and PutOnPPE. A state is entered by a command request for an action, with the state’s end condition defining when the action should terminate. The state performs the interactions necessary with the game to initiate the action and receive notice that the action has been accomplished. For the reactive agent design, all command requests originate from the player. However, as more sophisticated agents are developed for *Ground Truth*, agents can request these states to perform actions themselves, masking the communication between the agent and game world from the agent developer.

In creating our reactive agent, we constructed a hierarchical state machine (HSM) [2] by building meta-states that would chain together action states for more complex behaviors. For example, a player may request an agent evacuate a building on the other side of the game world. To create this behavior, one can construct a finite state machine (FSM), PlanEvacuation, which connects the MoveTo and Evacuate states for performing all actions needed to evacuate the distant location. PlanEvacuation serves as a child FSM to the master FSM, IdleState, which awaits commands from player and calls the proper child FSM to execute the command (see Figure 3).

III. RELATED WORK

A. Role of Emotion in Decision-Making

Since disaster response operations involve “in extremis” decision making, intense emotions associated with this type of environment can impair decisions by disrupting critical thinking. In fact, one’s emotions are a critical mediator in the types of decisions that are made and how people will ultimately behave. Emotion states serve as action tendencies that provide additional information to make judgments; especially when conditions are uncertain—as found in high-stress and high-consequence situations such as man-made and natural disasters. For example, [3] found that emotion states accounted for 34% of the variance in choices made when their expected utility was not obvious. Thus, representing the role of emotion in arriving at a decision may prove useful in assessing how people will ultimately behave.




Figure 3: Example of Master FSM calling Child FSM to handle evacuation and barricade commands

Converging lines of research suggest that a person’s attitude (which is a general opinion towards a person, object, or concept) influences behavior. A general theory supporting this research, the Theory of Planned Behavior, proposes that behaviors are influenced by 1) attitudes towards a specific behavior, 2) the subjective norms associated with acting out that behavior, and 3) the perception that this behavior is within a person’s control. This forms an action intention state, which then typically drives that person’s actual behavior [4], [5], [6].

A person’s emotional state often plays a large role in determining the ultimate behavior of the individual. According to the research by [7] and others, certain experiences may create general negative affects (such as the fight or flight impulse when a threat is perceived), which then may stimulate associations linked to fear and anger. How people respond can be a result of both goal- or moral-related decisions and their perceived emotion state. As mentioned above, assessments of the environment and the potential outcome can also temper behaviors. Consequently, angry persons might refrain from aggressive behavior if it would conflict with their goals or moral values. Accordingly, they may choose other behaviors that more closely align with their goals or values [8].

B. Modeling Emotion for Game-based Training Environments

Emotion modeling for NPCs is not a new concept. For commercial games, clever character development written

into game designs allow NPCs to reflect emotional growth and behavioral changes based upon the player’s choices [9]. Advancements in animation have allowed these NPCs to visibly express a wide range of emotions, giving a wider range of empathetic responses for the player. Yet, these techniques focus on just creating the illusion of emotions felt by NPCs, rather than attempting to understand and computationally model the human emotional responses within these characters. With commercial game AI techniques being viewed as an insufficient model for virtual characters in game-based training, or serious gaming [10], the necessity for more realistic human modeling has become a prevalent research area.

As noted in the prior section, emotions tie into a person’s motivation for deciding what actions to pursue. To computationally represent the role of emotion in motivation, researchers have found mapping perceptions to pre-defined emotional states as an effective method for virtual characters in these environments. [11] created agents for a derivation of the Iterative Prisoners Dilemma that mapped previous interactions with other agents to emotional states for determining whether to cooperate or defect. [12] incorporated a similar approach into a 3D predator/prey type game to aid a prey agent in reaching checkpoints without being detected by the predator.

Though several research efforts on emotion modeling exist, few instances of this research have transferred into game-based training environments. An explanation for this is the challenge in both modeling how emotion interacts with reasoning along with physical behavior. The Institute for Creative Technologies (ICT) [13] has undertaken a comprehensive effort toward including virtual humans into training systems, combining cognitive modeling with natural language generation and animation for representing to the trainee how emotion impacts these virtual characters. Since the *Ground Truth* platform does not represent virtual teammates with the same fidelity as the environments used by the ICT, the focus of our research has been on constructing psychology-plausible models of emotion at the agent decision-making level. This paper will later describe our first attempts at having *Ground Truth* communicate the emotions of the virtual teammates to the trainee.

Other researchers involved in modeling human cognitive processes have also investigated how emotion can construct a more realistic virtual teammate. The ICT virtual human research mentioned earlier uses the SOAR architecture [14] in forming their cognitive models. The SOAR community has made several contributions in modeling how emotion interacts with working memory, perception, and expression for NPCs [15], [16]. While SOAR tags perceptions with factors such as arousal, pain, and pleasure to make emotion an emergent property of the model, our research serves as a complimentary effort by directly defining the role of emotions in the decision making process.

C. Related Sandia Cognitive Modeling Work

Our work builds on research that initially sought a computational cognitive architecture that supports Human Natu-

realistic Decision Making [17] while incorporating “organic” factors such as emotion [18], [19]. The emphasis on emotions from the very beginning distinguishes this program of research from other cognitive architectures such as ACT-R [20] and SOAR [14]. Preliminary work attempted to combine a psychological model representing knowledge and cognitive processes with a physiologically-inspired model that provided the basis for incorporating organic factors. Using subsequently-developed simulation software that extended practicality, [21] developed a prototype human augmentation system based on discrepancy detection with respect to a task-based runtime cognitive model.




Figure 4. Conceptual cognitive architecture.

A high-level, psychological framework was fleshed out (Figure 4) to enable application within an embodied agent, as distinguished from a disembodied “decider” [22], [23]. At the same time, the Sandia Cognitive Runtime Engine with Active Memory (SCREAM) was developed to provide a practical cognitive simulation capability that supported that psychological framework. Emphasis on cognitive elements with activation-levels updated according to dynamics distinguishes SCREAM and most other Sandia cognitive simulation work from more common production-rule-based approaches.

The Sandia Human Embodiment and Representation Cognitive Architecture (SHERCA) is being developed to extend the psychological framework for modeling the behavior of humans as embodied agents while encouraging physiological plausibility. Conceptually, SHERCA fleshes out the high-level architecture down into subsystem models of perceptual memory, spatial memory, and action generation. SHERCA also refines the psychological model of decision-making with respect to action selection, providing a more detailed model of how emotional processes are integrated with the application of semantic and contextual knowledge.

Applying SHERCA for the conceptual and psychological model structure and SCREAM for the computational cognitive engine, [24] used runtime cognitive models to control the behavior of cognitive characters in a virtual 3D environment in a prototype training application emphasizing cultural awareness [22]. Standard embodied-agent simula-

tion techniques implemented in Sandia’s Umbra simulation framework [25] model the characters’ ability to “perceive” their environment, other entities, and those entities’ various attributes and actions.

Several other efforts at Sandia are developing computational cognitive modeling technologies that have foreseeable application in gaming. An example with current capability is the Cognitive Foundry [26], which provides (a) a theory-of-cognition-agnostic framework for developing cognitive models for intelligent agents and (b) a set of tools for automatically populating such models and evaluating them using statistical, machine learning, and visualization techniques. An example with applicability on the horizon is the development of a computational, neuro-physiologically plausible model of episodic memory [27].

IV. MODELING EMOTION IN VIRTUAL TEAMMATES

A. SHERCA

SHERCA was designed to correspond to the theories and supporting research mentioned in Section III [24]. SHERCA allows for multiple cues, cognitive perceptions, goals, action intentions, etc., to concurrently have some degree of activation. In SHERCA’s model of decision-making, once a cognitive perception—an element of perceptual/situational awareness—has been activated by cues in the environment, it may trigger activation of specific, intermediate goals that are consistent with higher-level goals and other active cognitive perceptions. For example, one high-level goal might be to protect family, and another, to protect oneself. Intermediate goals help support the higher-level goals by breaking down the goals into discrete tasks. The overall emotional state mediates activation of action intentions from the intermediate goals. As a consequence, the intended actions are a product of both the intermediate goals and the current emotional state of the simulated human. This emotional state may change dynamically, for example from very low to very high levels of anger, if the perceptions change. Action intentions that are contradictory with respect to goals can become concurrently highly activated due to the influence of emotion. At the same time, cognitive perception is influenced by a hierarchy of higher-level goals/directives or moral states, as well as state within a behavior (e.g., current step in a procedure).

In SHERCA, an impression of culture can be generated by varying a simulated human’s emotional response to particular perceptions. Cultures also exhibit variations within their high-level and intermediate goals. As a result, their intended and actual behaviors will show cultural uniqueness. The result is a complex set of behaviors that have certain emergent properties common to a particular group.

While SHERCA is intended eventually to achieve human-plausibility at a detailed level as various cognitive component models improve [27], initial emphasis has been on applying it to virtual embodiment. In this role SHERCA has been instantiated as a human-representative computational model through which a cognitive character recognizes patterns of stimuli in the environment and responds to those

stimuli according to current contexts, goals, and emotions [24]. Our main focus has been endowing characters with SHERCA's model of decision-making to select behavior-level actions rather than model detailed procedures or low-level control. Instead, we make use of non-cognitive AI/gaming techniques or features built into the virtual character model and virtual environment to implement those capabilities.




Figure 5 Model of decision-making to select actions in SHERCA.

B. SCREAM

SHERCA models in *Ground Truth* NPCs use SCREAM for the computational cognitive engine. SCREAM's representations combine aspects of ANNs and fuzzy logic networks. For example, we use ANN computations to model intensities of fear and anger, and convert them to fuzzy sets that are more easily understood by a human cognitive model builder. We now summarize key components of SCREAM, describe the basic cognitive elements, and sketch out their computational updates.

1) Basic cognitive elements in SCREAM

A **concept** is the fundamental semantic element in SCREAM. For convenience, concepts have names, but SCREAM associates no meaning with those names. To function in environments with multiple entities (e.g., things, creatures, features, etc.) of a given type requires a mechanism to associate concept activations with specific entities. SCREAM takes the simple approach of endowing concepts with slots. A **concept instance** associates slots with entities. Concept instances are created as needed. Each has its own activation state and is uniquely identified by concept and a vector of entity identifiers, which are merely labels to enable convenient interaction with people and other non-SCREAM system components. For example, *chases {22, 31}* identifies an instance of the concept *chases* (i.e., entity #22 chases entity #31). Thus, a concept is similar to a fuzzy predicate, but we do not claim that SCREAM implements any logic.

Contexts can be defined as meaningful perceptual representations that are based on recognizable patterns of stimuli, as well as, consistent with situation models, schema and theme-based representations of events. Context activation is governed by pattern recognition applied to the activation states of concepts that are the cues for/against that context. For example, the concepts *bicycle*, *clown*, *elephant* and *pop-corn* might be cues for the context *Circus*. A **context instance** is related to a **context** similar to the way that a **concept instance** is related to a **concept**. In SCREAM, a concept whose raw (input) activation is driven by the contextual pattern recognition process is also called a context.

Because contextual knowledge can include behaviors, SCREAM includes a context-to-abstract-action module that applies specified patterns to expand context instances into schema instances that describe behavior at a high level.

2) Updating cognitive state

At game start-up, each cognitive model instance is loaded with model definition and parameterization files. These files declare concepts and contexts and define context recognition patterns, context-to-behavior expansion patterns, and associations between concepts and emotions.

The main runtime cognitive state representation can be viewed as a dynamically-structured activation network of concept, concept instance, context instance and emotional state nodes. The *concept instance driver*, *semantic association network*, *context recognizer*, *emotional processes* and *context-to-abstract-action* modules of a SCREAM-based agent share the responsibility for updating the network structure and node activation. For example, the context recognizer uses a maximal unification criterion with respect to concept instances in maintaining the set of context instances. SCREAM *episodic memory* and *spatial memory* modules are not currently used in *Ground Truth*. Node activation values are always non-negative.

The instantiation and raw activation of input-level concepts are controlled from outside SCREAM via a concept instance driver. For example, a model definition might include the following input-level concept declarations:

```
c morning {}
c afternoon {}
c noon {}
c eats {ID1 ID2}
```

If the agent's perception module saw an entity it labeled *Fred* nibbling on one labeled *L19*, then it could call the agent's concept instance driver to set the raw activation level of the concept instance *eats {Fred, L19}* to, say, 0.7.

The semantic association network module updates the gross (output) activation of concept instances. Gross activation of a concept instance is computed from raw activation, spreading priming from concepts specified in the model, and top-down priming from contexts for which the concept is a cue. Priming is specified by global spreading- and top-down-gain gains and directed pairwise weights. For example,

```
sa noon {eats 0.2}
```

specifies a spreading activation weight of 0.2 from the concept *noon* to the concept *eats*.

Priming is distributed from concept nodes to concept instances. Summation with raw activation at instance nodes is followed by application of an activation function and rise/decay model. Concept instance node output updates are individually scheduled at rates within the alpha (8-13Hz) range, with frequency rising with activation level.

SCREAM currently includes two types of context recognition patterns. The first type of pattern is a template describing how specific concept instances act as inputs to associated instances of a given context and how the weights will be applied to the input activation levels.

For example, we can use

```

S Breakfasts {ID} {
  Eats-meal {{ID} 1.0}
  Breakfast-time {{{ 1.0}
  morning {{{ 1.0}
  afternoon {{{ -10.0}
} {ii 1.0}

```

as the pattern for a context whose activation reflects the awareness of the model that an observed entity is having breakfast. If, say, the concept instance *Eats-meal {99}* is active, then it will be an input to the context instance *Breakfasts {99}* with weight 1.0. The “{ii 1.0}” specifies that an intrinsic inhibitive bias of 1.0 will be applied when computing the activation level of an instance of *Breakfasts*. If *Breakfasts* is also a concept, then the *Breakfasts {99}* context instance output is the input to the corresponding concept instance, subject to capacity limitation. (See Figure 6.)

If χ is an instance of the context X , then the immediate activation level of χ is expressed by

$$a_{\text{immed}}(\chi) = \max\left(\left(\sum_k w_k a(p_k)\right) - \beta^{(X)}, 0\right).$$

Here $a(p_k)$ and w_k are, respectively, the gross activation and weight coefficient of the k^{th} input concept instance with respect to χ , and $\beta^{(X)}$ is the intrinsic inhibitive bias of context X . The gross activation level of χ is computed from $a_{\text{immed}}(\chi)$ by applying an activation function followed by a normalized leaky integrator. It is important for the activation function to be non-linear, and some activation function choices that SCREAM provides include:

$$\mathcal{A}(x) = \begin{cases} x & \text{if } x \leq 1 \\ 1 + \tan^{-1}(x-1) & \text{otherwise} \end{cases}; \text{ and}$$

$$\mathcal{A}(x) = \min(x, 1).$$

The latter is more intuitive for building models by hand, while the former appears better suited for machine learning.

Gross (output) activation levels of all context instances are updated at 5Hz (i.e., in the theta band). Outputs are made available synchronously to enable efficient implementation of simple capacity limitation.

A second type of context pattern is used for templates of context instances whose immediate activation is based on the activations all instances of the same concept that have matching values at specified slots. This pattern type has been useful in abstracting away slot values from concept/context instances and in approximately expressing a minimum quantity. For example,

```

XQ Eats-meal {ID1} {eats {{ID1 ID2} 1.0}} {ii 2}
XQ Breakfast-time {} {Breakfasts {{ID} 1.0}} {ii .5}.

```

In the first pattern specifies that currently active instances of the concept *eats* with matching first slots will contribute activation to the same instance of context *Eats-meal*. The “plain English” interpretation is that somebody is eating a meal if (s)he is eating at least three things. Immediate and gross activation of instances of contexts of the “XQ” type are computed similarly to the “S” type context activation.

This type of context pattern offers tremendous benefit in representation capacity for large numbers of concept instances that can be related in this fashion.

For example, subject to simple parameter choices, if we consider the model examples presented earlier in this section and activate the concept instances *eats {99 19}*, *eats {99 31}*, *eats {99 31}*, *eats {86 6}*, *eats {86 31}*, *eats {86 20}*, and *morning {}* all with raw activation 1.0, then the concept and context instance network will have the structure shown in Figure 6. The ability to define a set of context patterns that can give rise to recurrent network structures enables the models to be stateful even without additional memory components.




Figure 6: SCREAM runtime model with recurrent structure. Ovals represent concept instance nodes, and rectangles, context instance nodes; flow is from left to right except where indicated by arrows

A basic capability for modeling emotional processes in cognition [3], [24] has been implemented in SCREAM. SCREAM updates the level of activation of each emotion based on concept activation levels and parameters that specify how the concepts influence emotional state. Each concept can be associated with a level of activation and a weight coefficient for each emotion. For example, in an emotion parameters file,

```

cee clown 2 fear 0.6 0.7 anger 1.5 0.9

```

specifies that concept *clown* influences two emotions. For *fear* it has a weight coefficient of 0.6 and a target activation of 0.7, and for *anger* it has a weight coefficient of 1.5 and a target activation of 0.9.

The emotional processes module takes as input the overall activation levels of all concepts. The immediate activation level a_μ of emotion μ can be expressed as

$$a_\mu = \frac{\sum_{i=0}^{n-1} \sum_{j=0}^{m(i)-1} w_{\mu,i} \alpha(p_{i,j}) \xi_{\mu,i}}{\sum_{i=0}^{n-1} \sum_{j=0}^{m(i)-1} w_{\mu,i} \alpha(p_{i,j})}.$$

Here, $\alpha(p_{i,j})$ is the activation of the j^{th} of $m(i)$ active instances of concept i ; $w_{\mu,i}$ and $\xi_{\mu,i}$ are the weight coeffi-

cient and target activation for concept i associated with emotion μ ; and there are n concepts. Thus, the emotional state is computed similarly to center of mass.

The activation levels of the modeled emotions are converted into fuzzy-set representations whose membership levels determine the activation of corresponding concepts. For example, the level of *anger* can be represented by membership levels in the fuzzy sets *anger::Low*, *anger::Medium*, and *anger::High* that cover the range of *anger* activation levels. We currently use normalized Gaussian fuzzy set representations. Thus converted into concept instance representations, emotional state affects the activation of concept instances and context instances via the mechanisms described previously in this section.

C. Integration of SHERCA into Ground Truth Agents

We now describe architectural aspects of how SHERCA is integrated into *Ground Truth* agents. Each cognitive agent has a SHERCA-based cognitive model, which it mainly uses to continuously determine and update its high-level behavior. Each cognitive agent is an instance of the *ScreamAgent* class and has its own *SCREAM CognitiveModelObject* instance. The *ScreamAgent* class builds on the class of state-machine-driven agents described in Section III.C., and a cognitive agent uses this inherited capability to carry out the cognitively-selected behavior.

When a cognitive agent is created, it loads cognitive model definition (data) files whose contents include

- the concepts, contexts, and context patterns;
- context-instance to behavior/action conversion patterns; and
- spreading activation (priming) and emotional association parameters.

Currently, individual cognitive model definitions in *Ground Truth* differ only in emotional association parameters and levels of activation of high-level goals, reflecting differences in personality, culture and values. Setting activation of other specific concepts appropriately allows us to customize the generic model for each specific type of *Ground Truth* NPC.

To explain how the update of agent cognitive state fits into the game state update, we first view the latter loop at very high level:

- The respective Managers for fog, evacuation, traffic, etc., update the states of game elements not directly controlled by the agents.
- The Agent Manager updates the states of the NPC agents.

The AgentManager updates the states of the agents in an update cycle by doing the following:

- It updates the agents' physical states, based on their current states and game state external to the agents.
- It has the agents update their perceptual states, with help from the Perception Manager.
- It has the agents make decisions.
- It has the agents act on their decisions.

The last item is carried out by the cognitive agents' state-machine elements. Some actions take the form of making requests to the Game State, so that it can maintain consistency, instead of updating state directly.

To make decisions, a cognitive agent:

- Updates raw activation levels of input-level concept instances
- Iterates its internal *CognitiveModelObject* to the current game time.
- Passes the highest-ranked behavior option to its state machine component for execution.

The receipt of a user command that includes at least one argument, such as a location, corresponds to a concept that has a matching number of slots. Each unique argument value is translated into a symbolic label for activation of a concept instance. This label is subsequently converted back when the start state of a behavior is initialized.

In both the state-machine-based agents and cognitive agents, decision-making considers perceptual data, the most recent command received from the player, and various state data. For a state-machine-based agent, user commands directly set the high-level behavior that the agent will perform. A new user command results in the state machine popping states until one is reached that can dispatch the desired behavior. However, in a cognitive agent, the cognitive model determines the high-level behavior.

States in the state machine of a cognitive agent can access its emotional state for use in modeling affect within a behavior. For example, dialogue output takes into account emotional state. Generally, determination of low-level behavior, such as path planning, also makes use of separate algorithms that are called from within states.

D. Implementation of SHERCA Driven Teammate

We have developed a SHERCA cognitive model for *Ground Truth* NPCs and use specific model instances to govern their behaviors independently. The model currently contains roughly 80 concepts and contexts and over 50 context patterns. It is intended only to be sufficient for those NPCs, and we view its development as an exploration into the use of emotional cognitive models to increase the realism of the effects of player decisions.

In our initial cognitive model development and integration spiral, we constructed a minimalist cognitive model whose behavior selection enabled the cognitive NPCs to act like the state-machine-based NPCs. To begin our more recent spiral, we identified game situations for individual NPCs that we believed should evoke emotional responses. We then expanded the cognitive model to follow the SHERCA framework of decision making, and we identified situations in which it would be intuitive to a player when an NPC chose not to obey the most recent command it received, based on its emotional state and activation levels of its high-level goals.

Developing the minimalist cognitive model enabled us to exercise model elements needed for the behavior selection role. We defined input-level concepts for communicating:

- simple perceptive state, such as current location, noticing the toxic fog, or current action/behavior;
- receiving commands;
- endogenous state;
- declarative state knowledge, such as whether agent has PPE to wear, and whether it is being worn.

Note that in this discussion, names of input-level concepts will begin with a lower-case letter.

Because it is located between user-issued commands and the behavior-executing state machine, the cognitive model must keep track of the current command to (possibly) be acted upon. We defined contexts for achieving this capability. For example, we have a context whose instances model whether the current command is to move to a given location:

```
S Curr-cmd-move-to {LOC} {
  rcv-move-command {{LOC} 1.0}
  Curr-cmd-move-to {{LOC} 1.0}
  self-at {{LOC} -5.0}
  Task-some-other-loc-cmd {{LOC} -10.0} }.
```

Activation of a *rcv-move-command* instance results in activation of a corresponding *Curr-cmd-move-to* context instance. Because the activation duration of an input-level concept instances that model receiving a command is limited by a timer to approximate, e.g., conversation duration, a *Curr-cmd-move-to* context instance is self-stimulating once activated. (Recall that the activation function will modulate output activation.) The *self-at* negative cue models completing command execution or ignoring the command if the agent is already at the destination. The *Task-some-other-loc-cmd* negative cue, whose activation is also (context) pattern-driven, enables activation to be canceled by more recent reception of another command. Thus, the context pattern defines a recurrent relationship for sustaining self-activation until overcome by terminating cues.

SCREAM computes emotional state with respect to concept activation. In addition to basic recognition of certain stimuli or aspects of those stimuli, there are three particularly interesting categories of concepts that influence emotional state in *Ground Truth* agents:

- physical state or sensation;
- assessment of a situation;
- assessment of a situation including behavior of another agent.

The first category is currently limited to (the intuitively named) *feel-sick* and *feel-dying* concepts in *Ground Truth* ScreamAgents, which convert their health levels to activation levels of these concepts. The latter two categories fall into the Perceptual Awareness part of the SHERCA decision making model. An example from the second category is the concept that models awareness that the agent is endangered by the presence of toxic fog because (s)he is not wearing PPE. Examples from the third category correspond to being aware of carrying out an order that will result in people dying and being aware of carrying out an evacuation when the situation is appropriate. *Ground Truth* cognitive models include a permanently activated dummy concept and associated emotional parameters to define a base emotional state

and to provide resistance to mood swings due to low-activation concept instances.

Ground Truth cognitive models include three concepts that model the high-level goals of staying alive, maintaining discipline, and saving lives. Activation patterns of these concepts and those representing perceptual awareness give rise to the activation of intermediate goals and action intentions. Intermediate-goal-action-intentions (IGAI) that might imply a non-local behavior generally require the agent to receive a user command, such as the order to evacuate a specific block of the city, in order to be recognized. The high-level goal of maintaining discipline can help overcome local observations that might not apply at a task destination. IGAI whose recognition only requires local information can become activated without the NPC receiving a prompting command from the user, and even in opposition to a user command. An example of this latter type of IGAI is to seek to stop an evacuation in the presence of toxic fog to avoid hurting civilians.



Figure 7: Heads-up display for an agent after SHERCA integration. The green bar shows the agent's health level, yellow shows fear, and red shows anger level.

Emotion further contributes to selection among behaviors that are responses to the same situation. Higher levels of fear help prompt the behavior of putting on PPE in the presence of toxic fog without a command or approval from the IC, while sufficient levels of fear and anger in the same situation will cause the agent to panic and flee instead of communicating with the IC if the goal of maintaining discipline is not sufficiently high.

E. Discussion and Ongoing/Future Development

With a basic capability achieved, there are many directions for further development and research. Our plan is to improve the feedback to the player with respect to the NPCs emotional state by outputting context sensitive dialog and sound effects. By doing this we would seek to eliminate the need for the emotional status bars. However, we foresee this to be a labor-intensive process due to the necessity for a wide variety of relevant, yet different, verbal responses. In addition, we plan on semantically tagging areas of the game

world. This will provide our cognitive agents with richer perceptions of the game world and allow for more substantive cognitive models. Lastly, we plan on conducting further experiments to quantitatively demonstrate that our training objectives were fulfilled with the inclusion of cognitive NPCs. These tests will include subject matter experts from the training community.

Experience developing *Ground Truth* cognitive NPCs leads us to consider several directions for improvements to SCREAM, SHERCA, and integration of cognitive models into NPCs. We have noticed that emotion levels fall more quickly than what is intuitively appropriate. Implementing support for individual decay rates for concepts, and possibly computed emotion levels that they tend to evoke, could be a solution. We would like to understand how to appropriately model suppression or maintenance of a behavior for a duration that is most easily understood in terms of a particular length of time. We also need to consider whether/how to enable SCREAM to model emotions being directed at particular entities. In more complex environments, a model of attention would be needed. Finally, in the long term, SCREAM/SHERCA requires an understanding of how episodic memory is used in decision-making, with and without context learning.

Future scenario development from this work should look at events that span multiple cities, shared resources, and multiple attacks. High-level decisions from those scenarios impact response by adding extra constraints responders must operate in, thus affecting their emotional pressures. This opens up a new avenue for training even higher-level decision makers. Related developments could include training for leadership and teaming abilities. Another option is to adopt emotional models for the civilian population. For example, telling a region to evacuate while the toxic cloud is right over them would cause panic. Would the civilians continue to respectfully obey the first responders? This would enhance the training by providing support for civil disobedience.

V. CONCLUSION

Based on our early experiences with integrating SHERCA-driven cognitive models into *Ground Truth* NPCs we have seen positive results with increasing the realism of the training experience. The SHERCA cognitive model structure and methodology enables us to use SCREAM's combination of neural network representation and AI-based model description to effectively build psychologically plausible models that select behaviors as intended. We were able to incorporate fear and anger in an intuitive manner within SHERCA's modeling guidelines. Also, we are able to give the NPCs individual variations by specifying different emotional sub-models and setting different high-level goal activations at runtime. Within expected limitations, the cognitive NPCs make decisions and exhibit emotional states that are mutually consistent with the perceptions they are provided about the game state via its interface to the cognitive models. Thus, including emotion in NPCs via

SHERCA-based cognitive models increases the realism of effects based upon an IC trainee's decisions when playing *Ground Truth*.

VI. ACKNOWLEDGEMENTS

This work was funded through an internal research and development program at Sandia National Laboratories. Sandia is a multi-program laboratory operated by the Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

The authors would also like to thank John Linebarger, Paul Nielan, and Ann Speed for their comments and suggestions. Thanks also to Jason Honda, Edward Baker, and David Ko for their work contributing to this effort.

VII. REFERENCES

- [1] M. Wooldridge, *An Introduction to Multiagent Systems*: John Wiley & Sons, LTD., 2002.
- [2] M. Yannakakis, "Hierarchical State Machines," in *Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics, Lecture Notes on Computer Science 1872*: Springer-Verlag, 2000, p. 15.
- [3] M. Bernard and B. Smith, "The Effects of Emotional States and Traits on Risky Decision-Making," *Sandia National Laboratories Internal Technical Report (SAND 2006-7812)*, 2006.
- [4] I. Ajzen and T. J. Madden, "Prediction of Goal-Directed Behavior: Attitudes Intentions, and Perceived Behavioral Control," *Journal of Experimental Social Psychology*, p. 17, 1986.
- [5] M. Fishbein and M. Stasson, "The Role of Desires, Self-Predictions, and Perceived Control in the Prediction of Training Session Attendance," *Journal of Applied Social Psychology*, p. 25, 1990.
- [6] T. J. Madden, P. S. Ellen, and I. Ajzen, "A Comparison of the Theory of Planned Behavior and the Theory of Reasoned Action," *Personality and Social Psychology Bulletin*, p. 6, 1992.
- [7] L. Berkowitz, "Pain and Aggression: Some Findings and Implications," *Motivation and Emotion*, p. 16, 1993.
- [8] A. Bandura, C. Barbaranelli, V. Caprara, and C. Pastorelli, "Mechanisms of Moral Disengagement in the Exercise of Moral Agency," *Journal of Personality and Social Psychology*, p. 10, 1996.
- [9] M. Krawczyk and J. Novak, "Game Story & Character Development," *Thomson Delmar Learning*, 2006.
- [10] M. V. Lent, "Entertainment game AI vs. Serious game AI," in *Proceedings of the 27th SOAR Workshop*, 2007.
- [11] D. Chaplin and A. E. Rhalibi, "IPD for emotional NPC Societies in Games," in *Proceedings of the Affective Computational Entities Symposium*, 2004, pp. 51-60.
- [12] T. Hussain and G. Vidaver, "Flexible and Purposeful NPC Behaviors using Real-Time Genetic Control," *IEEE Congress on Evolutionary Computation*, p. 5, 2006.
- [13] P. Kenney, A. Hartholt, J. Gratch, W. Swartout, D. Traum, S. Marsella, and D. Piepol, "Building Interactive Virtual Humans for Training Environments," in *Proceedings of the 2007 Interservice/Industry, Training, Simulation, and Education Conference*, 2007.
- [14] P. Rosenbloom, J. Laird, and A. Newell, "The Soar Papers: Research on Integrated Intelligence," *MIT Press*, 1993.
- [15] A. Henninger, R. Jones, and E. Chown, "Behaviors that Emerge from Emotion and Cognition: Implementation and Evaluation of a Symbolic-Connectionist Architecture," in *Proceedings of the 2003 International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003, pp. 321-328.
- [16] R. Marinier III and J. Laird, "Toward a Comprehensive Computational Model of Emotions and Feelings," in *Proceedings of the*

- Sixth International Conference on Cognitive Modeling*, 2004, pp. 172-178.
- [17] G. Klein, "The Recognition-Primed Decision Model (RPD): Looking Back, Looking Forward," in *Naturalistic Decision Making*: Lawrence Erlbaum Associates, Inc., 1997, p. 7.
- [18] C. Forsythe and C. A. Wenner, "Surety of Human Elements of High Consequence Systems: An Organic Model," in *Proceedings of the IEA 2000 / HFES 2000 Congress*, 2000, pp. 3-839 - 3-842.
- [19] C. Forsythe and P. Xavier, "Human Emulation: Progress Toward Realistic Synthetic Human Agents," in *Proceedings of the 11th Conference on Computer-Generated Forces and Behavior Representation*, 2002, pp. 257-266.
- [20] J. R. Anderson and C. Lebiere, *The atomic Components of Thought*: Lawrence Erlbaum Associates, Inc., 1998.
- [21] C. Forsythe, M. Bernard, P. Xavier, R. Abbott, A. Speed, and N. Brannon, "Using Psychologically Plausible Operator Cognitive Models to Enhance Operator Performance," in *Human Factors and Ergonomics Society 47th Annual Meeting* Denver, CO, 2003.
- [22] M. Bernard, P. Xavier, P. Wolfenberger, D. Hart, R. Waymire, and M. Glickman, "Psychologically Plausible Cognitive Models for Simulating Interactive Human Behaviors," in *Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting*, 2005.
- [23] C. Forsythe, P. Xavier, and A. Speed, "Dynamic Model of Human Context Recognition and Understanding," in *Cognitive Systems: Human Cognitive Models in System Design* Santa Fe, NM, 2004.
- [24] M. Bernard, M. Glickman, S. Verzi, D. Hart, P. Xavier, and P. Wolfenberger, "Simulating Human Behavior for National Security Human Interactions," *Sandia National Laboratories Internal Technical Report (SAND 2006-7812)*, 2007.
- [25] E. J. Gottlieb, M. J. McDonald, F. J. Opper, J. B. Rigdon, and P. G. Xavier, "The Umbra Simulation Framework as Applied to Building HLA Federates," in *Proceedings of the 2002 Winter Simulation Conference*, San Diego, CA, 2002, pp. 981-989.
- [26] J. Basilico, Z. Benz, and K. R. Dixon, "The Cognitive Foundry: a Flexible Platform for Intelligent Agent Modeling," in *Proceedings of the 17th Conference on Behavior Representation in Modeling and Simulation*, Providence, RI, 2008.
- [27] S. Verzi, S. Taylor, T. Caudell, M. Bernard, and D. Morrow, "An Adaptive Resonance Theory based Computational Model of the Hippocampus," *in preparation*, 2008.