

Modeling Rules of Engagement in Computer Generated Forces

Rick Evertsz

AOS Group

Level 1, 594 Elizabeth Street, Melbourne, Victoria 3000, Australia
rick.evertsz@agent-software.com

Frank E. Ritter

College of Information Sciences and Technology

The Pennsylvania State University, University Park, PA 16802

frank.ritter@psu.edu

Simon Russell

QinetiQ Ltd.

Cody Technology Park, Ively Road, Farnborough, Hants, GU14 0LX, UK
sprussell@qinetiq.com

David Shepherdson

AOS Group

Wellington House, East Road, Cambridge, CB1 1BH, UK

david.shepherdson@agent-software.co.uk

Keywords:

Modeling and Simulation, ROE, Agents.

ABSTRACT: *Rules of Engagement (ROE) are driven by a mix of legal, military, and political factors. These dimensions can interact and overlap in subtle ways and must be carefully crafted to be easy to apply in combat situations without jeopardizing mission outcome and the warfighter's right to self-defense. Although trial and error may have sufficed in the past, the growing complexity of conflicts and the military and political ramifications of ineffective ROE (e.g., a friendly fire incident), make a simulation-based ROE evaluation system a high priority. This paper describes ROE3, a human behavior-modeling tool that supports tactics-independent representation of ROE. In our approach, ROE are defined as meta-knowledge that act as a constraint on the tactical choices selected by the synthetic entity. This is key to the flexibility of the system — tactics and ROE can be freely mixed and matched to investigate their interactions.*

1. Introduction

Rules of Engagement (ROE) are:

Directives issued by competent military authority that delineate the circumstances and limitations under which United States forces will initiate and/or continue combat engagement with other forces encountered (JOINT PUBLICATION 1-02, 1999).

A mix of legal, military, and political considerations drives ROE. These dimensions can interact and overlap in subtle ways particularly where multinational forces are concerned and where conflict arises with an unconventional foe (e.g., an out-of-uniform enemy in

urban terrain). Such ROE need to be carefully crafted to be easy to apply in combat conditions without jeopardizing mission outcome and the soldier's right to self-defense. The political, legal, and tactical complexity of the modern military environment makes it difficult to evaluate ROE, leading to undesirable incidents in the field. Consequently there is a pressing requirement for simulation tools that allow ROE to be developed and evaluated in advance of a conflict.

Prior to the research reported here, the modeling of ROE in a simulation context has been fairly rudimentary (e.g., setting fire permissions and target priorities in the OneSAF Testbed Baseline) or non-existent. Kalus and

Hirst (1998) mention the possibility of adding ROE to Soar models, but concrete results were not reported.

A key problem in modeling how military personnel behave during a conflict is to understand how ROE are interpreted with respect to the tactical environment and mission goals to be achieved. Most current technology does not model the following factors:

- ROE,
- how a person selects a tactic within the scope of several ROE being applied simultaneously, and
- how this selection process is affected by the type of operation, psychological aspects of the person, and moderating influences experienced by that person (such as stress or fatigue).

Whether during drafting, training or conflict, ROE present a variety of challenges and tradeoffs. For example, ROE designed to limit escalation can put friendly forces at risk, particularly where the enemy has no formal ROE. Evaluating draft ROE is a major challenge in the absence of a conflict arena.

Although trial and error may have sufficed in the past as a development process, the growing complexity of current and future conflicts, and the military and political ramifications of ineffective ROE (e.g., a friendly fire incident), make a simulation-based ROE evaluation system a high priority. With the support of an ROE simulation tool, problematic ROE could be detected in advance of a conflict by testing draft ROE against a library of definitive scenarios. Our implementation of this approach, ROE3 (Rules of Engagement Evaluation Environment), offers a number of significant benefits:

- Candidate ROE can be systematically evaluated for effectiveness and unforeseen implications.
- Use of ROE3 will encourage standardization of ROE formalization, thus reducing the cost and time taken to evaluate changes to ROE.
- ROE3 includes modeling human variability factors like stress and fatigue, thereby supporting the evaluation of how these moderators can affect ROE handling in a wide variety of circumstances.

2. ROE Background

ROE are drafted to comply with legal, political, and military considerations. Effective ROE are defined as constraints on action; they do not deal with specific tactical situations, but delineate the tactical solutions that are permissible. Following on from some brief

background material on ROE, this section describes our *software perspective* on ROE.

The most common classification of ROE is in terms of the political/legal/military dimension. This classification scheme highlights the motivation for the ROE. An example of a legal constraint would be to prohibit attacks on places of worship or on hospitals. A military constraint might be to require sighting of the target before subjecting it to indirect fire, and a political constraint could prohibit entry into civilian dwellings (to avoid alienating the civilian population and compromising the political goals of the mission).

Whether political, legal, or military, the drafting of ROE varies depending on whether there is a state of war or OOTW (Operations Other Than War). In the US, Standing Rules of Engagement (SROE) are the core ROE for peace time operations (other than those involving assistance to domestic authorities in handling civil disturbances). SROE permit engagement for the purposes of self-defense, whereas wartime ROE permit US forces to fire upon enemy targets that do not present an immediate threat. Wartime ROE are permissive, which means that warfighters may engage the enemy even if the enemy is not currently adopting a threatening posture. In contrast, SROE are typically restrictive — opposing elements can only be engaged if they are clearly behaving in a hostile manner.

The range of possible ROE is very large. However, as a conflict unfolds, the mission will drive the ROE in a particular direction, for example, from conduct-based to status-based rules in response to the emergence of hostile activity. Conduct-based ROE are defined in terms of the behavior (conduct) of the enemy, whereas status-based ROE are contingent upon the state of the situation faced. Conduct-based ROE are rules that describe the steps to be followed when enemy hostile intent can be inferred from their behavior. On the other hand, status-based ROE dictate a sequence of actions to be taken when the situation appears to endanger US forces. For example, if a US soldier on patrol guarding a unit encounters an armed person trying to infiltrate the perimeter, the ROE specify what can be done based on the situation, without inferring intent.

2.1 Software Perspective on ROE

The legal/political/military classification of ROE highlights the motivation for particular ROE. However, it is too high-level to serve as a practical classification scheme, thus FM 27-100 (The Judge Advocate General's School, U.S. Army, 2000b) further categorizes ROE into ten groups that address the more practical, military aspects of ROE.

The legal/political/military and FM 27-100 ROE classification schemes are important in drafting ROE and for training. However, the goal of ROE3 is to implement ROE in synthetic entities. From this perspective, we need to ascertain whether there are classes of ROE that will require different software implementation techniques. This section examines ROE from this software perspective and is based solely on the ROE Handbook (The Judge Advocate General's School, U.S. Army, 2000a). The five classes of ROE described below are implemented in ROE3.

2.2.1 The Right to Self-Defense

A key feature of all ROE is that they never override the right to self-defense. No matter how restrictive the ROE, US forces are able to bypass any restriction that endangers their lives. The inherent right to self-defense is a rule with special status in ROE3. It acts as a caveat on all other ROE.

2.2.2 Restrictive ROE

Many ROE prohibit action unless specified conditions apply, for example: "Do not engage unless enemy exhibits hostile intent or action". Such ROE act as a restraint on action and are termed Restrictive ROE. SROE are often restrictive in nature.

2.2.3 Permissive ROE

In contrast to Restrictive ROE, some ROE define behavior that is permissible under particular conditions and are termed Permissive ROE. Such ROE enable the individual to bypass the process of checking that a proposed action does not violate Restrictive ROE, and thereby facilitate rapid response in a fast-changing situation. Because the hostile intent of the enemy has been established, wartime ROE are largely permissive in nature.

2.2.4 Preference ROE

Some ROE define preferred actions, for example: "Armed force is a last resort". Given a number of alternative courses of action that all achieve the same short-term military goal, ROE can be applied to select the one that best serves the overall goals of the mission. The "Armed force is a last resort" rule often applies in peacekeeping operations, where the political goal is to move the host nation towards peace. Unlike Restrictive ROE, such rules do not explicitly exclude courses of action, but allow the individual to select the one that best satisfies the political/military goals.

2.2.5 Criterion-Defining ROE

Criterion-Defining ROE are included to disambiguate ROE that are criterion based. To illustrate, many ROE refer to **Hostile Intent** as a criterion for the engagement of potentially hostile elements.

3. ROE3 Architecture

The architecture of ROE3 is outlined in Figure 1. Each behavior agent in ROE3 consists of an ROE Framework that interacts with the CoJACK overlay and the JACK agent architecture. The Integration Layer translates messages between the agent and the synthetic environment (SE). This translation process can be quite complex, for example, aggregating raw SAF (Semi-Autonomous Forces) percepts into a higher-level percept such as "Entity ID0023 is pointing his rifle at me".

The ROE and tactics layers embody the agent's cognitive and skilled action capability, that is, what it knows and what it knows how to do. This capability is constrained by the agent's cognitive architecture to fall within the range that can be supported by the human cognitive system (e.g., limits on working memory). The moderator layer forms a feedback loop with the cognitive architecture, for example, further limiting working memory capacity in times of stress. The level of stress itself, in turn, can be affected by the agent's assessment of the situation. Moderators can be based on internal factors such as the level of adrenaline as well as external factors like temperature and humidity.

Although the ROE Framework is central to the modeling of ROE, the behavioral realism is underpinned by the JACK and CoJACK layers. We will now briefly outline these two layers before describing the ROE Framework in more detail.

3.1 JACK

JACK is a mature, cross-platform environment for building, running and integrating multi-agent autonomous systems. It is built on a sound logical foundation: BDI (Beliefs, Desires, Intentions). BDI is an intuitive and powerful abstraction that allows developers to manage the complexity of the problem. It is based on work by Bratman (1987) on situated rational agents. In JACK, agents are defined in terms of their beliefs (what they know and what they know how to do), their desires (what goals they would like to achieve), and their intentions (the goals they are currently committed to achieving).

JACK supports meta-level reasoning using the same representation (**plans**) as it does for tactical reasoning. Meta-level plans provide essential support for implementing the interaction between ROE and tactics.

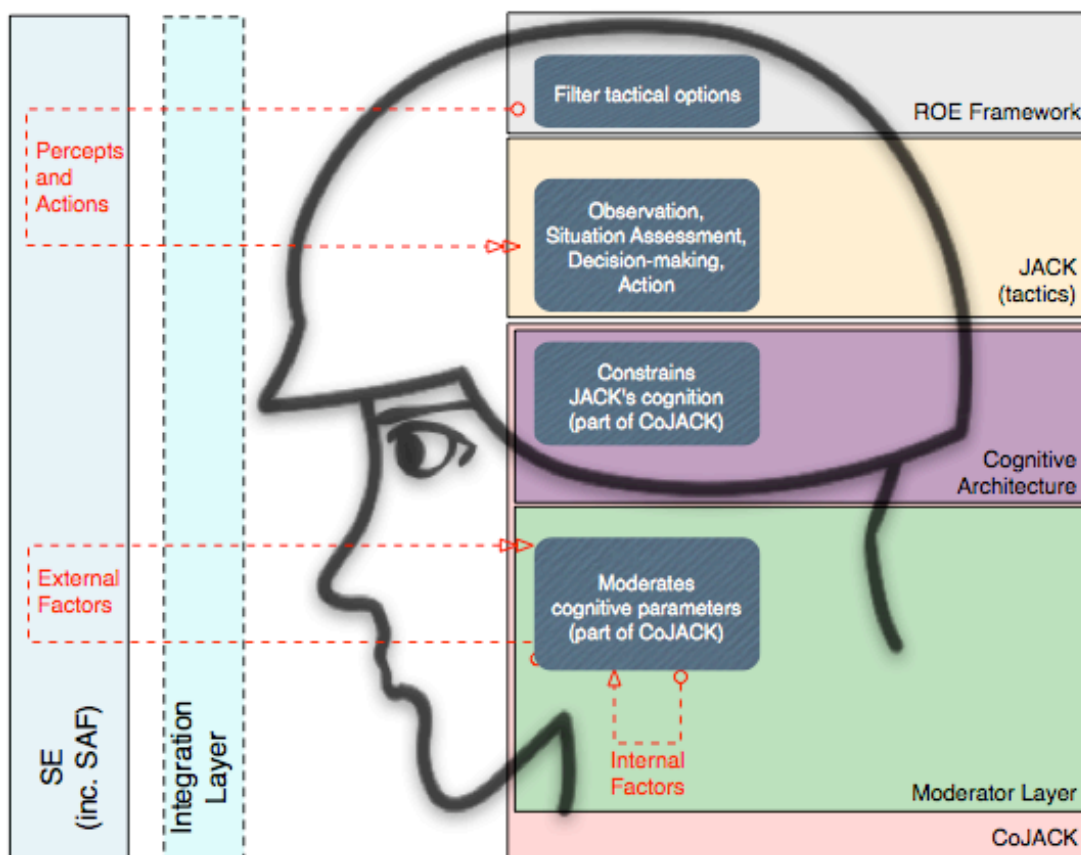


Figure 1. ROE3 Architecture

JACK has a number of features that make it well-suited to modeling human cognition in the context of Computer Generated Forces (CGFs), including (a) autonomous agents, (b) socially structured agent teams, (c) a reactive, pattern-directed inference mechanism, (d) a notion of commitment to a Course of Action (CoA), and (e) a user-friendly graphical plan representation.

3.2 CoJACK

In military simulation, relatively little attention has been given to modeling human cognition, and how this affects human behavior. Although languages like JACK are well suited to modeling rational reasoning, they are not designed to address behavioral influences that result from moderating factors. Examples of moderating factors are internal ones like emotion and fatigue, and external ones such as temperature and suppressive fire (Hudlicka & Billingsley, 1998; Silverman et al., 2000; Ritter and Norling, 2006). The Improved Human Behaviour Representation (IHBR) project (MoD project: RT/COM/3/006) addresses this shortcoming by extending

JACK so that it can take human variability and behavior moderating influences into account. In addition, it augments JACK with a set of constraints that mirror key properties of the human cognitive system so that behavior variation can be simulated in a principled manner. This extension is **CoJACK**.

3.3 Integration Layer

The Integration Layer consists of the CGF Integration Layer (CGF-IL) and Babel Box. It provides the mechanism to allow CoJACK to task entities within a SAF system. The messaging mechanism is bi-directional: as well as being able to receive tasks from CoJACK, the SAF can send situation awareness data via the Babel Box back into the ROE3 agent.

The CGF-IL uses a set of generic messages based on a general ontology of behaviors, tasking, and situation awareness data. These messages are used to extract data from and send commands to the Babel Box. The Babel Box is configured to operate with a particular SE and

converts the generic message into a SAF-specific action. This means that the CGF-IL would not need to be changed if another SAF were used. The particular SAF used in this experiment was the OneSAF Testbed Baseline (OTB) v2.0 International.

3.4 OneSAF Testbed Baseline (OTB)

The ModSAF family provides a flexible and scalable synthetic battlespace capability. Entity behaviors, tasks, data, interfaces, etc., are all provided in reader files and source code. This means that the ModSAF family of CGF systems is relatively simple to modify, and therefore was the preferred choice for this study.

The candidate systems for this project were OTB v2.0 and JointSAF (JSAF), where JSAF was the preferred system in the original proposal. JSAF is a tri-service simulation system, including detailed modeling of air, land, and sea assets. However, JSAF lacks a comprehensive dismounted infantry model. The final scenario, set in an urban environment with Individual Combatant (IC) interactions, would require that extra functionality be added to JSAF to fulfill the needs of the project.

OTB is an intermediate system and is intended as a stepping-stone to the next generation of ModSAF, namely the OneSAF Objective System (OOS). OTB primarily caters for land-based behaviors and interactions, however it does include the representation of air and sea assets. In addition, OTB includes specialized Dismounted Infantry SAF (DISAF) extensions and behaviors developed to support US Army simulations. The existence of IC modeling and behavior made OTB the preferred choice for the SE in this study.

3.4.3 Sensor Models

OTB models visual and aural sensors as well as electronic sensors such as radar. A SAF IC entity will visually detect and classify targets based on their range, aspect and the overall visual environment (i.e., lighting, obscuration). There is also an aural model that will detect footsteps, unobserved fire, and fall of shot.

3.4.4 Internal Models

OTB models the individual as an upper and a lower body. The sensor model and weapons are mounted on the upper body, and the lower body hosts the entity's mobility model. This simplification of the human form allows the SAF to simulate in real time multiple entities on standard computer equipment. However, for the purposes of this experiment, the OTB model lacks the ability to simulate posture in any great depth. SAF postures are limited to standing, crouching, kneeling, sitting, and prone. There

are also weapon states: shouldered, deployed (lowered), and deployed (raised).

There is currently no weapon representation other than the three states above. The OTB model does not allow the entity to point its weapon in any particular direction. However, it is possible to alter the internal appearance of an entity to extend the existing postures and weapon states. In addition, OTB has an internal targeting model that can be used to derive the weapon direction based on the entity's current selection as its best target. Extending these models in this area will be important for improving the fidelity of ROE simulations because knowing posture and direction are important for inferring intent. This study has made some steps in this direction.

3.4.5 Tasking Entities

Entities in OTB use a hierarchical approach to behaviors and tasking. Each behavior library can control a single entity or a group of entities by manipulating lower-level behaviors. The most basic behaviors will interact directly with the physical model and the simulated world.

3.5 Modifications to OTB to support ROE

The SAF system was designed to model physical and behavioral aspects of an entity and collective-level tasking. This includes physical motion, interactions between entities and low-level behaviors. The introduction of external control via the Babel Box required some modification to the way that individual entities are tasked.

In normal operation, the SAF manages all low-level behaviors and interactions with other entities and the physical environment. The user interface allows entities to be tasked, by a human controller, with high-level behaviors, such as travel on a route, make a hasty attack, defend a position, and so on. These discrete tasks are followed by the entity in sequence.

On testing the close-coupled control of an entity via the Babel Box, it was found that the speed at which the SAF could switch tasks is limited to around 2 to 5 seconds. This is not fast enough to allow for the rapid task changes needed to control an individual combatant entity. Therefore, a modification was required in order to improve the response of the SAF to external control. The modifications allowed an entity to be controlled via a single task object — removing the need to switch task-frames over the network, which is a time-intensive process. By wrapping all the basic behaviors into a single task object, it was possible to tightly control the execution of orders being received via the Babel Box. This control was necessary to ensure that the SAF responded in a predictable manner. The behaviors were encapsulated in a

single additional task library. This provided direct control of low-level IC behaviors, including:

- Movement to a location (x,y coordinate)
- Orientation while moving and halted (degrees from map North)
- Speed (kph)
- Posture stance (standing, crouched, prone, sitting)
- Weapon stance (stowed, deployed, raised)
- Special postures (waving gun)
- Fire at a target (specific target callsign)
- Fire at a location (x,y location plus a z-offset)
- Throw grenade/rock at a location
- Mount/dismount a vehicle
- Communication across a radio net (emulates vocalization)

4. ROE Framework

The ROE Framework combines a means for implementing ROE together with a methodology for representing ROE and tactics in a way that enables the ROE to affect the chosen CoA. In this framework, ROE are represented as meta-knowledge, that is, knowledge that is used to reason about knowledge. Specifically, ROE are represented using meta-plans that make a choice amongst the plans that are applicable to the current situation (plans are used to represent tactics in JACK). The methodology requires that tactics be annotated with their ROE-relevant effects. For example, if a tactic involves shooting an entity, and there are ROE that deal with lethal force, then that tactic must be labeled as one that applies lethal force. Only then can the meta-plans determine whether an ROE that prohibits lethal force should exclude a tactic that involves shooting an entity.

Figure 1 shows the ROE Framework, CoJACK, and the JACK agent as separate layers in the ROE3 architecture. Although each layer has a distinct responsibility, any synthetic entity (JACK ROE agent) will be made up of all three layers if it is provided with ROE, behavior moderation, and tactical reasoning capabilities. Because the three layers reside within the one agent, they can be given access to one another. For example, the ROE Framework can access the situation assessment of the tactical layer (JACK) and can be influenced by the moderators in CoJACK. This is what one would expect

from a psychologically plausible architecture, for example, ROE processing should be affected by the level of fear represented in the CoJACK layer and should also be able to make use of the tactical layer's assessment of the level of threat in the current situation.

4.1 ROE Methodology

ROE meta-plans act as restraints/preferences on tactical plans. This requires that the meta-plans have a means of determining the ROE-relevant effects that a given tactical plan has on the environment. There are three aspects to this:

1. If a plan performs an ROE-relevant action, then it is labeled as such. Typically, the majority of plans in a synthetic entity do not execute ROE-relevant actions but perform functions such as processing percepts, assessing the situation, and implementing the adopted course of action. Labeling plans that perform ROE-relevant actions ensures that the ROE meta-plans do not waste effort processing plans that have nothing to do with ROE (this purely an efficiency measure).

2. Plans that are ROE-relevant must be annotated with a list of the actions that are relevant to ROE, for example, if a plan performs a **Shoot entity X** action.

3. Because ROE do not normally use concrete terms like shoot, a mapping must be defined between abstract ROE actions and the concrete ones used in tactical plans. For example, rather than refer to actions such as shoot, ROE will tend to refer to lethal force.

The current implementation of ROE3 not only requires the model builder to manually annotate tactics with their ROE-relevant actions, but necessitates that a given tactic be annotated with any ROE-relevant actions that necessarily follow its adoption. This is because a CoA is normally represented as a sequence of plan fragments (in the interests of sharing those fragments across CoAs). For example, consider a Platoon Assault tactical plan where the Commander starts monitoring the location of the Assault Platoon, instructs the Fire Support Platoon to commence firing, waits for the Assault Platoon to arrive in the engagement area, commands the Fire Support platoon to cease firing, waits for confirmation that the Fire Support Platoon is quiescent, and tells the Assault Platoon to commence the attack. Strictly speaking, instructing the Fire Support Platoon to commence firing does not constitute the application of lethal force. ROE3 needs to be able to determine, however, that this is an action that indirectly results in lethal force.

There is a range of possible approaches to determining if an action in a given plan, X, necessarily results in an ROE-relevant action in a plan invoked by plan X. The

most straightforward solution is to require that the model builder explicitly annotate such plans, in other words, annotate plan X so that it is clear that it results in the application of lethal force. This is the approach adopted in ROE3 on the basis that when coding the tactical plan, the model builder knows that it will invoke another plan that actually applies the lethal force. The other extreme is to infer this through an automated analysis of the tactics library. Though technically feasible, this automated approach lay outside the scope of this research effort. Nevertheless, the next section briefly describes how such an analysis could be automated.

4.2 Automated Analysis of ROE Dependencies

To automatically analyze ROE dependencies, it is necessary to determine whether a given plan *entails* an ROE-relevant action in some other plan. This is a classical search problem (Nilsson, 1980) that asks the question: Is there a path from the current state to another state in which an ROE-relevant action is performed, and more critically, is that the only path open to the synthetic entity?

This is non-trivial because synthetic entities operate in a dynamic environment where it is not always possible to predict how other entities in the simulation will behave. Realistically, the adversary always has a range of options (e.g., surrendering), and so there is limited value in attempting to factor the adversary's options into the automated analysis of ROE dependencies. Nevertheless, it is possible to envision cases where it is important to reason about the ROE implications of an adversary's potential responses. These cases relate to CoA generation and analysis.

CoA generation and analysis is the process of determining how to achieve the mission goal. This involves predicting and limiting the tactical responses open to the enemy, and relates to ROE insofar as the ROE constrain the permissible options. Thus, it is necessary to perform some form of look-ahead search when trying to determine what implications adopting a plan has for ROE. The search space will generally be larger in cases where the adversary's response must be factored into the search. This is an area where computer support will help understand the implications, as simulation here can be more complete than unassisted human consideration.

Synthetic entities normally take concrete values as inputs (simulated percepts) and produce other concrete values as

outputs (actions to be performed on the SE). However, when searching a state space of possibilities it is impractical to use concrete values because the space will grow exponentially and will become effectively infinite. To overcome this problem it is necessary to use uninstantiated input variables. Here the goal is to characterize the general behavior of the simulation model, that is, compute the abstract actions that occur in response to abstract inputs. This process of abstractly executing a computer program (which is effectively a simulation model) is termed Abstract Interpretation (Cousot and Cousot, 1977).

Abstract Interpretation is a commonly employed solution to the problems of automatic program analysis. The general idea is to glean information about a program by running it on abstract specifications of data objects, rather than the data objects themselves. This can be performed during the simulation run or offline (in advance). The search space will be significantly more tractable if it is performed at runtime, rather than offline. This is because many of the input variables will be partially instantiated, effectively reducing the size of the search space further.

4.3 ROE Meta-Plan Filtering

Figure 2 illustrates the behavioral loop adopted for models in ROE3. This is effectively Boyd's OODA (Observe, Orient, Decide, Act) loop (Coram, 2002) and characterizes the combat decision cycle in terms of four phases. The behavior model acquires perceptual input from the environment, assesses the input, decides what actions to perform, and finally performs the selected action(s).

1. In the Observe phase the synthetic entity acquires the environmental changes through its sensors.
2. These percepts are then processed in the Orient phase where the entity assesses the situation and draws conclusions about what is happening.
3. In the Decide phase the entity collects the CoAs that apply to the situation, as assessed in the Orient phase. The ROE Framework is then used to filter out and prioritize those CoAs that are impacted by ROE.
4. Finally, in the Act phase the entity performs the action(s) in the selected CoA.

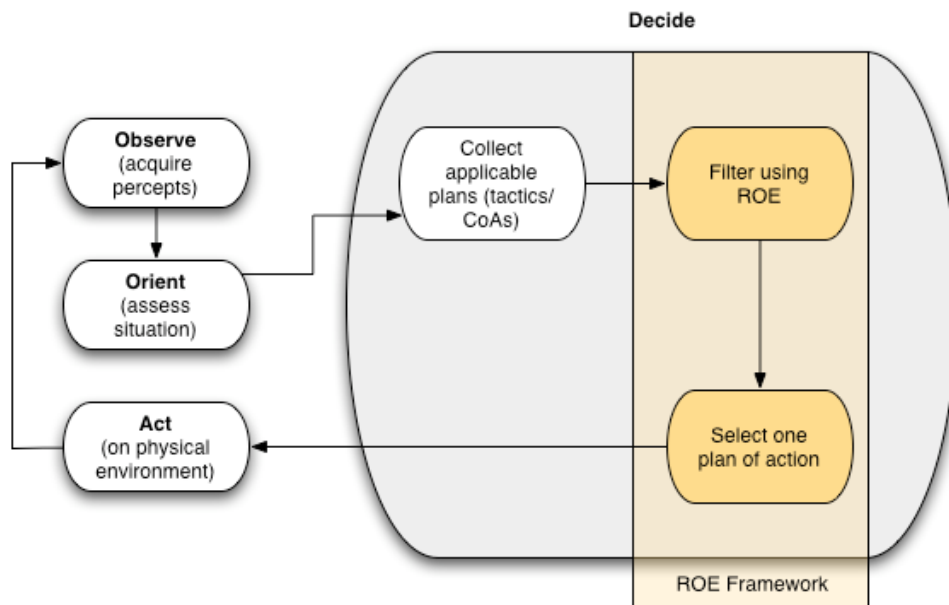


Figure 2. OODA Loop incorporating ROE

When they are first defined, the ROE are classified into their appropriate categories (i.e., Self Defense, Permissive, Restrictive, and Preference). At runtime, when using the meta-plans to filter tactical options, ROE3 goes through the following steps:

1. First check whether any ROE sanction the candidate tactical plan on the basis of the Right to Self Defense. If any do, then accept the candidate plan; there is no need to check the Permissive or Restrictive ROE.
2. If the Right to Self Defense is not relevant, determine if Permissive ROE sanction the candidate plan. If any Permissive ROE apply to the candidate plan in the current situation, accept the candidate plan — there is no need to check Restrictive ROE. This reflects the semantics of Permissive ROE and fulfills their intended purpose — in the interests of rapid response in time-critical situations, bypass the application of Restrictive ROE if certain key conditions are satisfied.
3. If neither the Right to Self Defense nor Permissive ROE are applicable, check the candidate plan using Restrictive ROE. If any Restrictive ROE match the candidate plan in the current situation, then the candidate plan is rejected.

Each of the above three steps can invoke sub-tasks to help determine the applicability of the ROE to the current situation. For example, when checking if the Right to Self

Defense sanctions the candidate tactical plan, determine which actor, action, and targets the candidate plan is responding to, in other words which entity or entities the candidate plan is proposing to apply lethal force to, and for those entities, what action each entity is performing and what targets that entity is applying that action to. This information is then used in to determine whether the entity has demonstrated Hostile Intent (by accessing Definitional ROE plans that specify what constitutes Hostile Intent).

5. Scenario

ROE3 has been evaluated using OTB v2.0 on a scenario with a number of variations adapted from the ROE Handbook (The Judge Advocate General's School, U.S. Army, 2000a). The scenario involves a Marine convoy approaching a rebel-manned roadblock. There are three main scenario alternatives (m1 to m3), with another three minor variations that derive from the impact of the fear moderator. In variation m1, the rebel leader aims his pistol at one of the Marines. In variation m2, a rebel soldier throws a rock at the lead HMMWV, and in variation m3 the rebels rush the Marines at weapons ready. In addition, in all three variations the Rebel Leader waves his weapon on first contact with the Marines. The Commander agent's response to this is solely dependent on his assessment of how hostile this action is, and leads to the one of the three minor variations in the Commander's response, noted below. This assessment

derives in part from the Commander agent's fear moderator.

5.1 Marine Agent Design

The scenario comprises five types of Marine agent: Commander, Interpreter, Marine Soldier, Radio Operator and Driver. The Marine agents have a Situation Assessment capability and an Action capability. The Situation Assessment capability comprises sub-capabilities including Perception, World Modeling, and Threat Assessment. The Action capability similarly comprises sub-capabilities like Walk, Adopt Threatening Posture, and Speak. Each of these capabilities can be specialized in a given agent type, for example, the Commander agent is given a more sophisticated Situation Assessment sub-capability.

5.2 Rebel Agent Design

There are two types of Rebel agent: Rebel Leader and Rebel Soldier. The overall Rebel agent design is similar to that of the Marine agent, but the Rebel agent has a different range of capabilities (e.g., Withdraw, Throw Rock).

5.3 Moderated Behavior – Fear (Task Appraisal)

A secondary goal of this study was to demonstrate moderated synthetic behavior, and in particular, moderated ROE handling. To this end, a "fear" moderator was implemented in the demonstration scenario.

Providing a validated model of fear is constrained by the following three factors:

1. The definition of fear – fear might exist on multiple levels, ranging from loss of life, down to anxiety about the success of a particular task or plan. For this scenario, fear is defined as an emotional aversion to harm from the adversary, and this influences behavior to avoid the things that cause this feeling. At low levels this change might not lead to changes in strategies, but changes in how existing strategies are implemented, for example, actions might be performed more slowly. At higher levels of fear, different strategies might be executed, such as not firing, or firing when not appropriate.
2. Cues – fear arises out of interaction with the environment, and the cues or stimuli to create fear must be available. The cues necessary for becoming fearful are included in the Marine behavior models.
3. Access to data to validate fear – because of its nature, fear can be difficult to study in laboratory settings, and the ability to find it in naturalistic settings

has its own problems. A related literature that we can draw on is the task appraisal literature. This literature, for example Lazarus and Folkman (1984), notes that people make appraisals of how difficult a task will be, and how they will be able to cope with it, including what resources they bring to the task. These are combined to create an appraisal of the situation. Where sufficient resources are available to cope with the demands of the task, the task can be viewed as challenging. When the task requires more coping resources than are available, the task can be perceived to be threatening. When these appraisals are manipulated and knowledge held constant, performance is seen to vary. Stress responses in physiology are also observed based on differences in this process (e.g., Kirschbaum et al., 1993).

We thus examined the physiology of fear and then the related concept of task appraisal. Existing models of fear, physiology, and task appraisal were examined for lessons in implementing the existing model.

There have been a few cognitive architectures that implement how fear can influence behavior, including:

1. PSI is a cognitive architecture designed to integrate cognitive processes, emotions, and motivation (Bartl and Dörner, 1998). The architecture includes six motives (needs for energy, water, pain avoidance, affiliation, certainty, and competence). These motive/emotional states and their processes modulate cognition.
2. Barry Silverman (2004) has created an architecture called PMFServ. This architecture uses reservoirs of various resources to provide for different affective moderators and mental states.

Developing a validated model of fear was outside the scope of this study. The fear moderator was added to demonstrate that the CoJACK architecture facilitates investigation of the interaction between moderators and ROE.

In the scenario, the level of fear moderates the threat assessment (in the Orient phase of the OODA loop), which in turn influences the chosen course of action. During the application of the ROE, the level of fear moderates the perception of Hostile Intent. If Hostile Intent is perceived, Self Defense can be used as a justification for the application of lethal force. Each agent can be specialized to assess the threat of events in different ways. For example, when the Marine agents perceive the Rebel Leader agent waving its weapon, the increment in fear level will vary depending on each agent's previous experience of such situations. If such events have had lethal consequences, the level of induced

fear will be higher than if such events always ended up being a bluff.

5.4 Evaluation of Scenario Runs

Due to the nature of the study, the evaluation was qualitative in nature. In each scenario variation, the Marine agents applied their ROE to select from the available tactical options, leading to alternate outcomes. Their performance was no longer uniform across all three scenarios.

5.4.1 The Effects of ROE

The Commander agent's primary goal is to overcome the roadblock, and three initial tactical options are available:

1. Demand that the rebels allow the convoy to pass.
2. Forcefully demand that the rebels move the roadblock.
3. Fire upon the rebels.

The Restrictive ROE forbidding lethal force excludes the third option while the Preference ROE that favor minimum force will select the first option. The Commander agent will try the second option if the Rebel Leader refuses to comply with the initial less-forceful demand.

5.4.2 The Effects of Behavior Moderation on ROE

If the Rebel Leader agent waves its weapon, the Commander agent's response is solely dependent on its assessment of how hostile this action is. This assessment derives in part from the Commander agent's fear moderator. If the ROE are fixed, the fear moderator is the key factor in the variability of the Commander agent's response. The level of fear is determined by the Commander agent's previous experience of observing weapon waving in an adversarial context. The fear level is lowest if the agent has had similar experience with no negative consequences, and highest if the experience has involved wounding or lethal results for friendly forces or elements under US protection.

In the model, the level of fear affects the hostility threshold, that is, the point at which weapon waving is assessed as hostile. If the weapon waving is assessed as hostile, ROE permit the use of lethal force as a response.

If the weapon waving is not assessed as hostile (e.g., because the Commander agent's fear level is low), the Commander has a number of response options — from verbally defusing the situation, to lethal force. The Commander agent's ROE reject lethal force because

hostile intent has not been observed. ROE related to Minimum Force then are used to choose to verbally defuse the situation rather than selecting more forceful non-lethal responses.

5.4.3 SAF Issues

Although ROE3 was successfully tested and demonstrated using OTB, OTB is not an ideal SAF for investigating ROE in urban scenarios. One of the problems is that it can be difficult to control individual combatants in OTB because they have their own inbuilt behaviors. Another problem is that the visualization is not detailed enough for many types of ROE investigation. A 3D visualization environment would be beneficial in areas such as OOTW and urban operations. Working with a lightweight simulation like dTank (Ritter et al., in press) would make the development and testing of ROE faster, but would represent a tradeoff of accuracy in the simulation.

6. Discussion

The primary goal of this research program was to develop a tool that supports the investigation of the interaction between ROE and tactics. The key design goal was to separate the representation of ROE from that of the tactics. It would have been trivial, and not very useful, to have hard-wired the ROE into the tactics themselves. Had the ROE been hard-wired into the tactics, it would have been impractical to alter the ROE in order to investigate their effects on tactics. JACK's support for encapsulation allows generic behavior fragments to be grouped into collections that can be inherited by behavior model instances, allowing reuse of ROE and tactics components.

The primary goal has been achieved by representing ROE as meta-knowledge, that is, knowledge that is used to reason about knowledge. Specifically, ROE are represented using meta-plans that make a choice amongst the plans (tactics) that are applicable to the current situation. Plans are used to reason about and act upon the state of the SE, whereas meta-plans are used to reason about the reasoning itself. They are fully-fledged plans with the added power of being applied independently at the point that the agent makes a choice between the available tactical options. This means that meta-plans can make use of the full power of JACK plans and are not handicapped in any way.

The other major advantage of this approach is that meta-plans have full access to the internal state of the agent and so can make use of conclusions drawn by the situation assessment plans. Furthermore, because meta-plans have all the capabilities of normal plans, they can be moderated. This is a very powerful benefit of the approach because it allows the investigation of the interaction between

variables like fatigue and the ability to successfully apply ROE to the current situation.

6.1 ROE and Human Variability

This research also makes a contribution to the goal of incorporating realistic human variation in CGF models. In many current SEs, the individual entities in the simulation will execute the same task in the same way, ignoring differences between individuals, and even the variability of a given individual over time. In the real world, this is not the case. The choice of strategies and the ordering of sub-strategies will vary across individuals and will vary for a given individual across time. When such variance is not included in a model, it makes adversaries, allies, and neutral personnel too predictable because they will always do the same thing at the same time in the same way. Modeling this variance can be extremely important for sensitivity analysis and for training.

CoJACK allows JACK models to take account of the time humans take to reason, shortfalls in memory, and the effects of physical and emotional factors on cognition. This then allows the creation of CGF models that exhibit more psychologically-grounded behavior variation. A major strength of our approach is the use of meta-plans to represent ROE. This makes ROE-related reasoning subject to the same effects as the rest of the cognitive model. For example, if the working memory capacity of an agent were reduced as a result of short-term stress, one would expect situation awareness, tactical reasoning, and ROE handling to be adversely affected. CoJACK provides this benefit and continues to be developed under UK MOD's IHBR program (RT/COM/3/006).

6.2 ROE as Procedural/Declarative Knowledge

In ROE3, ROE are represented as meta-plans. Like normal JACK plans, meta-plans are a procedural representation. Arguably, procedural representations are appropriate when the knowledge is well rehearsed and in some sense *compiled*. Declarative representations are more suited to cases where the knowledge has to be interpreted. Being a procedural representation, JACK plans are not suited to tasks like analyzing the logical consistency of ROE. ROE meta-plans act as a *recipe* for applying ROE to the current situation, rather than as an abstract representation of the *meaning* of those ROE. ROE3 would be greatly enhanced by providing a declarative ROE representation that complements the current procedural one. This can be achieved in JACK by using **beliefs** to represent declarative ROE, and we look forward to extending ROE3 to use JACK belief structures for declarative ROE in tandem with the current meta-plan representation.

7. Acknowledgements

The development of this software was supported by AFRL/MLKH award FA8650-04-C-6440. This work was also supported by the UK MOD's Analysis, Experimentation and Simulation corporate research programme (Project No: RT/COM/3/006). We thank Bill Knarr for helpful comments on ROE, and Colin Sheppard for useful discussions about CoJACK.

8. References

- Bartl, C., & Dörner, D. (1998). PSI: A theory of the integration of cognition, emotion and motivation, in *Proceedings of the 2nd European Conference on Cognitive Modelling*, 66-73.
- Bratman, M. E. (1987). *Intention, plans, and practical reasoning*. Cambridge, MA (USA): Harvard University Press.
- Coram, R. (2002). *Boyd: The fighter pilot who changed the art of war*. New York, NY: Back Bay Books/Little, Brown and Company.
- Cousot, P., & Cousot, R. (1977). Abstract interpretation: A unified framework for static analysis of programs by construction or approximation of fixpoints, in *Proceedings of the Fourth Annual ACM Symposium on Principles of Programming Languages*, 238-52.
- Hudlicka, E., & Billingsley, J. (1998). Proceedings of the 7th Conference on Computer Generated Forces and Behavioral Representation, in *Representing behavior moderators in military human performance models*
- JOINT PUBLICATION 1-02 (1999). *DOD Dictionary of Military and Associated Terms*, 452.
- Kalus, A., & Hirst, A.J. (1998). Soar Agents for OOTW Mission Simulation, in *Command & Control Decision Making in Emerging Conflicts - 4th International Command & Control Research & Technology Symposium*
- Kirschbaum, C., Pirke, K. -, & Hellhammer, D. H. (1993). The Trier Social Stress Test—A tool for investigating psychobiological stress responses in a laboratory setting. *Neuropsychobiology*, 28, 76-81.
- Lazarus, R. S., & Folkman, S. (1984). *Stress, appraisal and coping*. New York: Springer Publishing.
- Nilsson, J. (1980). *Principles of Artificial Intelligence*. Palo Alto, CA: Tioga Publishing Company.
- Ritter, F. E., & Norling, E. (2006). Including human variability in a cognitive architecture to improve team

simulation. In R. Sun (Eds.), *Cognition and multi-agent interaction: From cognitive modeling to social simulation* (pp. 417-27). Cambridge, UK: Cambridge University Press.

Ritter, F.E., Kase, S.E., Bhandarkar, D., Lewis, B., & Cohen, M.A. (in press). dTank updated: Exploring moderator-influenced behavior in a light-weight synthetic environment, in *Proceedings of the 16th Conference on Behavior Representation in Modeling and Simulation*

Silverman, B. G. (2004). Human performance simulation. In J. W. Ness, D. R. Ritzer, & V. Tepe (Eds.), *The science and simulation of human performance* (pp. 469-98). Amsterdam: Elsevier.

Silverman, B.G., Johns, M., Shin, H., & Weaver, R. (2000). Performance moderator functions for human behavior modeling in military simulations, *Prepared for Human Behavior Program, Defense Modeling and Simulation Office*

The Judge Advocate General's School, U.S. Army (2000a). *Rules of Engagement (ROE) Handbook for Judge Advocates*. Charlottesville, VA: U.S. Army's Center for Law and Military Operations.

The Judge Advocate General's School, U.S. Army (2000b). *FM 27-100, Legal Support to Operations, Department of Defense Dictionary of Military and Associated Terms*. Washington DC: Headquarters, Department of the Army.

Author Biographies

RICK EVERTSZ is a cognitive scientist at AOS Group. He is currently working on the incorporation of cognitive and affective constraints into BDI agents.

FRANK RITTER is one of the founding faculty of the College of IST, an interdisciplinary academic unit at Penn State to study how people process information using technology. He works on the development, application, and methodology of cognitive models, particularly as applied to interfaces and emotions. He is an editorial board member of *Human Factors* and *AISB Journal*. His review (with others) on applying models in synthetic environments was published in 2003 as a HSIAC State Of the Art Report.

SIMON RUSSELL is a principal scientist at QinetiQ Ltd., specializing in modeling and simulation research. He has supported a number of projects using JSAF in Joint and non-Land applications, and has been working with the OneSAF TestBed in the UK since 2002.

DAVID SHEPHERDSON is a software engineer at AOS Group. He has worked on numerous agent-based simulation and decision support systems for both military and commercial applications in Australia, Europe, and the USA.