Requirements for Future SAFs: Beyond Tactical Realism

Robert E. Wray, PhD
Soar Technology, Inc.
Ann Arbor, MI
wray@soartech.com

Heather A. Priest, Melissa M. Walwanis, Katherine Kaste
NAWCTSD
Orlando, FL
{heather.priest, Melissa.Walwanis, katherine.kaste} @navy.mil

ABSTRACT

A key component of realistic and effective training in simulation is the behavior of semi-automated forces (SAFs). SAFs provide opponents, friendly forces, and other dynamic entities within the simulation. In most cases today, SAFs are designed and implemented to be tactically realistic; that is, they take actions that carry out good tactical decisions. As a result, SAFs are typically evaluated in terms of the realism or "fidelity" of their actions to the tactical situation and not with regard to training effectiveness.

We contend SAF tactical realism is a necessary but incomplete requirement for cost-effective and training-effective deployment of SAFs for simulation-based training. SAF behavior should also be modulated by scenario/exercise goals and also by the learning needs of individual trainees. In practice, these additional requirements tend to surface during delivery of training, requiring human instructor/operator teams to intervene. Interventions both increase the cost of simulation-based training and potentially lower the aggregate effectiveness of that training: delivering an appropriate experience at an apt time to the trainee is contingent on the attention and action of the instructional team. Further, as SAFs are increasingly used in mixed live-virtual-constructive training situations, SAFs that consider only tactical decisions will further limit scalability and increase the operational cost of LVC training.

In response, we suggest that imbuing the training system with the capability to understand and support scenario goals and individual training needs can make SAFs more practical for everyday training. We present examples of adaptation and variation that may be important for training but that are not typically embedded in a tactical SAF. We discuss the implications of these missed requirements and outline suggestions for incorporating interpretations of learning context in future simulation systems based on experience researching and developing such a capability. Finally, we outline methods for verifying and validating SAFs designed to meet these additional requirements.

ABOUT THE AUTHORS

ROBERT E. WRAY, PhD, is a senior scientist at Soar Technology. His research encompasses many areas of artificial intelligence including knowledge representation and ontology, agent-based systems and agent architectures, artificial intelligence in education and training, machine learning, and cognitive science.

Heather A. Priest is a Senior Research Psychologist at the Naval Air Warfare Center Training Systems Division (NAWC TSD). She received her M.S in Experimental Psychology from Mississippi State University and her Ph.D. in Human Factors Psychology from the University of Central Florida. At NAWC TSD, her research revolves around fidelity and realism of semi-automated forces (SAF), live, virtual and constructive training, and adaptive training for aviation. .

Melissa Walwanis is a Senior Research Psychologist and NAVAIR Associate Fellow at the Naval Air Warfare Center Training Systems Division (NAWCTSD). She leads the Training, Readiness, & Usability in System Technologies (TRUST) Laboratory which focuses on researching and acquiring effective training systems for multi-team instruction in LVC environments. She holds an M.S. in Industrial and Organizational Psychology from the University of Central Florida and is pursuing a Ph.D. in the same field at the Florida Institute of Technology.

Katherine Kaste is a Research Psychologist at NAWCTSD in Orlando, FL. She holds an M.S. in Human Factors and Systems from Embry-Riddle Aeronautical University (ERAU). Her research interests include cognitive psychology, systems design, and human-machine interface.

Requirements for Future SAFs: Beyond Tactical Realism

Robert E. Wray, PhD
Soar Technology, Inc.
Ann Arbor, MI
wray@soartech.com

Heather A. Priest, Melissa M. Walwanis, Katherine Kaste
NAWC TSD
Orlando, FL
{heather.priest, Melissa.Walwanis, katherine.kaste} @navy.mil

INTRODUCTION

Semi-automated forces (SAFs) play a key role in simulation-based training, providing opponents, friendly forces, and other dynamic entities. In mixed live-virtual-constructive (LVC) training environments, SAFs complement human trainees and role players with *constructed* roles, potentially interacting with human trainees in both virtual training (e.g., flight simulators) and live training (e.g., in-flight aircraft instrumented to the simulated environment) contexts. We use "SAF," rather than similar terms like *intelligent agents* or *human behavior models* because "semi-autonomous" highlights today's reality: ensuring synthetic entities behave in ways that support training objectives typically requires human oversight and intervention, a defining feature in the coinage of "SAF" (Ceranowicz, 1994).

The primary requirement for a SAF is that it behaves realistically in the context of its mission and the interactions it supports. For example, in beyond-visual-range (BVR), air-to-air (a2a) tactical training, threat SAFs need to fly realistically enough that they look similar on radar to what a live aggressor pilot might present. In this paper, we refer to such behavioral realism as (SAF) *behavior fidelity.* SAFs must generate sufficient behavior fidelity to support the targeted training without producing artifacts that introduce negative training transfer or gaming (Pew & Mavor, 1998). SAFs are often evaluated in terms of the fidelity of their actions to the tactical situation. Evaluation often compares expected or actual human behavior to SAF behavior (Gluck & Pew, 2005; Wallace & Laird, 2003).

The level of fidelity needed in an application is contingent on the training context. For example, BVR training requires a lower level of fidelity than within visual range (WVR) training: WVR SAFs will need to move tactically in three dimensions while the BVR model can realistically execute long-range engagements in (mostly) two dimensions. Understanding what behavioral fidelity is actually needed is a critical step in requirements analysis, as it is for other elements of the training system, such as the visual representation of the training environment. Researchers in virtual training have developed a *Layered Fidelity Framework* (LFF, Stacy et al., 2013) that organizes training-simulator fidelity requirements into a series of layers and maps the layered fidelity requirements onto trainee subjective experience. The LFF outlines a "Behavior Models" layer; behavior fidelity in this context describes behavior-modeling requirements for training. This paper explores how behavior fidelity interacts with trainee experience and what those interactions reveal about the requirements for SAF behavior fidelity.

The LFF anticipates the need for simulation and scenario manipulations designed to support learning. Going further, we contend that SAF behavior should be directly modulated by training goals and by the learning needs of individual trainees. Such *SAF Training Fidelity* is a component of behavior fidelity but is distinct from typical behavioral realism requirements. For example, a high fidelity tactical model may not be the best model for training across the training continuum. Similarly, tactically realistic (i.e., high fidelity) maneuvers may not ensure that the trainee gets the learning experience needed or required at a given time. We suggest SAF training fidelity is comparable to tactical realism in importance and it is critical for SAFs to be both training effective and cost-effective to deploy.

Below, we illustrate ways SAF training fidelity relates and interacts with typical realism requirements for SAFs and outline some operational gaps that result when SAFs with good realism lack the ability to consider training context. To meet the full scope of behavior fidelity requirements, we recommend imbuing the training system with the capability to understand training goals and individual training needs, presenting examples of variation that may be important for training but that are not typically embedded in a tactical SAF. The examples are taken from on-going work focused on enhancing an existing SAF tactical flight simulation system. Finally, we outline methods for verifying and validating SAFs designed to meet these additional requirements.

SAF TRAINING FIDELITY

As suggested in the introduction, we required behavior fidelity of a SAF behavior model is a function of both behavioral realism and SAF Training fidelity, or

Behavior Fidelity = f (behavior realism, SAF training fidelity) (1)

SAF training fidelity characterizes the degree to which the SAFs in a simulation-based training environment provide learners with individualized experiences that are targeted to training goals and adaptive to trainee capability. SAF training fidelity should also minimize the need for human intervention to support these capabilities. We expand on each element of this definition, using examples from generic, BVR fighter concepts (Shaw, 1985) to illustrate. Subsequent sections explore the functional interactions between behavior realism and SAF training fidelity and options for coordinating them to achieve the desired behavior fidelity for a given training application.

Targeting Training Goals

Simulation-based training experience is usually segmented into individual episodes of training (a "scenario"). In the typical case, each practice episode has a defined purpose, such as practicing some skill or demonstrating some level of competency relative to a training objective. In air-to-air tactical training, a scenario might involve conducting two intercepts against separate target groups, one right after the other. A schematic of the scenario is illustrated in Figure 1. In this case, the training objective are likely to be 1) executing the intercepts on timeline (achieving range/bearing relationships at particular times) and 2) employing weapons against the targets. Embedded in these training goals is the instructor-designed challenge of conducting two distinct but successive intercepts, so that as soon as the first timeline is completed, the next must be initiated. The behavioral realism required for the target groups is to fly an intercept against the trainee. When the scenario flows as designed, behavior realism is sufficient to deliver the required behavior fidelity. In Figure 1, assume there is sufficient range between the leading and trailing threats so that it is appropriate for the trainee to consider them as separate groups.

Figure 1. A "Successive Intercepts" Training Scenario.

Learners can introduce variation in the scenario that can threaten the goals of the presentation. Consider the situation shown in Figure 2. The trainee in this case has undertaken a defensive maneuver that impacts the timeline against the first threat and threatens the ability of the system to present a second, distinct intercept opportunity (a primary training goal). The training environment will dictate the specific recovery actions that are needed in this situation. If an instructor is present, the instructor might guide the trainee in making additional maneuvers to recover the timeline for the initial intercept. If an operator is present, the operator could intervene to turn the trailing threat as shown so that it remains "out of the battle" during the first intercept. This maneuver represents a poor tactical choice, but it may be the most appropriate solution for achieving the desired training goals given the trainee's action. For example, if this was a live or live/constructive training scenario, the cost of reset or even having the trainee maneuver to recover might not be justified. Manipulation of the target offers an obvious solution to allow the trainee to practice the training goals.

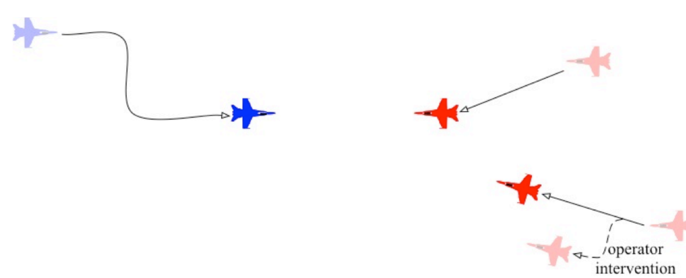

Figure 2. Trainee Action in a Successive Intercepts Training Scenario Can Impact the Required Behavior

Thus, even for relatively simple scenarios, human intervention may be required to ensure that the desired training situations are presented. The missing element is not tactical realism, but the ability of the system to direct SAF behavior to meet training goals; i.e., SAF training fidelity.

Adapting to Trainee Capability

As a trainee progresses in learning a complex task, such as the skill of flying a tactical aircraft, different knowledge, sub-skills, and abilities will be differentially exercised and developed. The trainee's attention, motivation and internalization of experience, among other factors, result in an individual, and sometimes possibly unique trajectory in this knowledge and skill space. Adapting to trainee capability allows a training environment to respond to these individual trajectories to improve overall learning outcomes, such as greater efficiency of time-on-task to reach competency or the accelerated development of emergent skills (Scales, 2006; Shute & Zapata-Rivera, 2012).

Consider again the training objectives and successive intercepts scenario illustrated in Figure 1. The level of challenge in this training setup derives from the initial geometry and the ability of the pilot to regard the threats independently. The initial geometry can be changed to provide different levels of challenge while still enabling the presentation of the desired training objective. In Figure 3a, the groups' initial headings and aspect to one another will make it more difficult for the trainee to separate and focus on the closer target group. In Figure 3b, he initial geometry makes it simpler to regard the two aircraft as independent and separable threats.

(a) A More Challenging Initial Geometry

(b) A Less Challenging Initial Geometry

Figure 3. Modifying the Level of Challenge.

Dynamic adaptations as the scenario executes can also be made sensitive to trainee proficiency. In the Figure 2 situation, the speed of the trailing SAF could be slowed once it was clear the trainee had failed to meet the employment timeline for the first threat (i.e., targeting of training goals). Whether speed adjustments were allowed could be influenced by the estimated proficiencies of the trainee.

In these examples, we have presented adaptation as a form of *scaffolding*; that is, assisting trainees in meeting practice challenges that are assumed to be beyond their current capability without assistance (Pea, 2004). However, adaptation based on estimated proficiency is not necessarily tied to assistance. For example, a range of different strategies may be used to prepare trainees for learning (Bransford & Schwartz, 1999) and in some cases presenting a situation beyond a trainee's ability may be used as a motivational technique. For example, a trainee that did not maintain separation with the trailing threat could be targeted and shot down by that threat during the intercept of the leading threat. Although the training effectiveness of this action can be debated, more salient for this analysis is that the training system can dynamically adapt simulation events in response to the estimated proficiency level. As we discuss below, one of the advantages of more automated approaches to proficiency-based adaptation is that the instructional strategy can be modified and tailored during execution, in addition to interventions that directly change the trainee experience (Wray & Woods, 2013).

Minimizing Human Intervention

SAFs today do not automatically provide more than behavior realism. However, as the examples above suggest, behavioral requirements for SAFs extend beyond providing a realistic tactical experience, either to reflect real-world situations or to reinforce instruction (e.g., consequences of an error). In this situation, the built in or "native" behavior fidelity is only a function of the tactical realism:

$$BehaviorFidelity_{native} = f\,(behavior\ realism) \tag{2}$$

Based on the prior discussion, the required behavior fidelity for the application may require behaviors that cannot be generated natively in this system. In this case, human intervention is required to achieve the require behavior fidelity. There are the significant limitations that arise when a system depends on human intervention to meet SAF training fidelity needs that are not natively addressed by the training system.

First, when variation is needed, the instructor/operator must recognize this need, choose a course of action, and then issue directives to individual SAFs. Alternatively, instructors may choose to script behaviors for future scenarios. In either case, manual control diverts instructor attention from more critical tasks such as mentoring and guidance, which are more important for learning (Chi et al., 2001). Additionally, manual control limits the span-of-control of the instructor in practice, making it less feasible to present the richest detail of the actual domain when such richness would be appropriate for the trainee's level of skill.

Second, understanding how to deliver the right kinds of instruction (or practice) at the right time is a challenge for highly trained instructional designers, educational psychologists, and teachers. Instructors and operators often may not have training of best practices in instruction. They may use rules of thumb that derive from tradition rather than evidence. Using a rule of thumb is both potentially inconsistent with best practices for instruction (e.g., "punish the mistake") and also results in inconsistent application across the training community.

Third, requiring that instructors and operators are present to meet SAF training fidelity requirements limits the contexts in which (efficient) training practice can occur. While there is likely to be little harm in training with and against SAFs with sufficiently high tactical fidelity, the impact of the practice on the learner is lessened. In the ideal case, the training system would encourage and support *deliberate practice* that encouraged a focus on the deliberate development of subskills and progression toward mastery (Ericsson, Krampe, & Tesch-Romer, 1993; Mayer, 2008). Instructional design of practice scenarios and a community of learners ("lessons learned") play as large or larger roles in supporting the development of expertise "outside of class." However, training systems that could autonomously and directly support programs of deliberate practice have the potential to significantly enhance the value and impact of simulation-based training.

In summary, SAF training fidelity requirements tend to surface during delivery of training and, as a consequence, they require instructor/operator teams to intervene. Intervention not only increases the cost of simulation-based training, it also lowers the aggregate effectiveness of that training because delivering an appropriate experience at an apt time to the trainee is contingent on the attention and action of the instructional team. Further, as SAFs are increasingly used in mixed live-virtual-constructive training situations, SAFs that only consider tactical decisions will further limit scalability and needlessly increase the operational cost of LVC training.

REQUIREMENTS FOR IMPROVED SAF TRAINING FIDELITY

Having described the concept of SAF training fidelity and motivations and examples of improved SAF training fidelity, we now introduce specific requirements for including capabilities that can enhance the native behavior fidelity of training systems. The roles of SAF training fidelity in the training system and the relationship of SAF training fidelity to tactical realism shape specific SAF behavioral and representational requirements.

Behavioral Requirements

SAF training fidelity and behavior realism can relate to one another in different ways and the framing of this relationship influences how requirements are articulated and measured. Consider, for example, the four categories of pilot behavior summarized in Table 1. Each class is summarized by adherence to doctrine (doctrinal), realism (something a human pilot could do), and the functional effectiveness of the behavior (a behavior that is appropriate in the context of mission goals). For this paper, we assume that all doctrinal behavior is functional, which reflects the purpose of doctrine in training. The "good" intercept path illustrated in Figure 1 is doctrinal, realistic, and functional. The path shown in Figure 2 is realistic (a trainee could perform this action) but not functional. There are also potential maneuvers that might result in some advantage in closing with the trailing threat before the leading threat (e.g., if the trailing threat was a faster aircraft with longer-range sensors and weapons). In this case, the resulting behavior might be realistic and functional but not doctrinal.

As these examples highlight, the first three classes of behavior can be applied to any role-player in the training environment, whether a human trainee (in either the live or virtual context), a role-player (including entities manually directed from an operator panel) or SAFs. The fourth class, unrealistic, is by definition not reflective of human behavior, but the realism judgment comes from the perspective of what the entity is doing relative to its mission. For example, a human operator who "warps" some entities into more desirable positions is generating "unrealistic" behavior in the context of the simulation.

Table 1. Different Categories of LVC Role-player Behavior.

Behavior	Doctrinal	Realistic	Functional	Description
Doctrinal	Y	Y	Y	Behavior that is consistent with doctrine. “Doctrine” implies both general tactical doctrine (such as how to conduct a maneuver) but also adherence to rules for the particular exercise or scenario (e.g., altitude restrictions for different mission states and roles).
Realistic & Effective	N	Y	Y	Behavior that is not consistent with doctrine but that reflects what a human role-player might do in a similar circumstance. This behavior does not conflict with the goals and constraints of the mission.
Realistic but ineffective	N	Y	N	Errors. Behavior that reflects what a human role-player might do in a similar circumstance. Non-doctrinal and not functional.
Unrealistic	N	N	Y N	Behavior that is non-consistent with what human role-players would or could do in similar circumstances. Unrealistic actions can be functional or non-functional.

Figure 4 illustrates, conceptually, how the categories relate to one another in a behavior space. In the diagram, its silhouette represents the current state of an entity. Each state affords a space of possible choices and actions available to the entity. As individual actions are taken, a behavior (a trajectory in the behavior space) is generated. The resulting behavior can be classified into the categories enumerated in Table 1. In reality, of course, these categories are not sharply distinct and judging whether some behavior is one category or another may be difficult and subjective. For this analysis, however, the categorization offers some useful distinctions for articulating SAF training fidelity requirements.

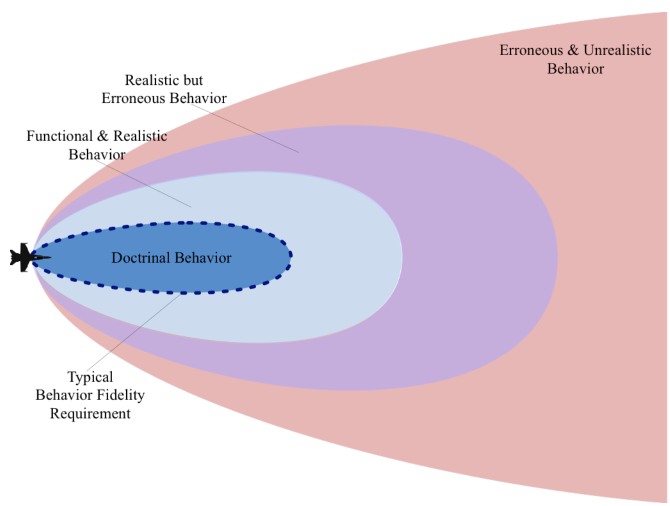

Figure 4. SAF Training Fidelity Requirements Depend on Behavioral Roles that SAFs Should Support for Training.

1. In the majority of today’s SAF systems, behavior fidelity requirements are usually focused on doctrinal behavior: SAF behavior needs to be realistic and consistent with the doctrine of the force being modeled. Similarly, the evaluation of SAFs often focuses on whether behaviors adhere to doctrinal prescriptions. The figure illustrates the behavior fidelity requirement and adherence to doctrine as equivalent sets because this is the norm. Limiting behavior fidelity to doctrinal behavior omits many behaviors likely to be encountered in the real world.

2. Regardless of the required behavior realism, SAF training fidelity will often require SAF behaviors that are not limited by doctrine. As the examples from above suggested, providing realistic responses in support of training goals may be important for training, even if those responses are not strictly doctrinal. This observation reinforces the earlier analysis of when and why intervention is needed. If the system requires for training that an entity exhibit some behavior that is outside of doctrine and doctrine alone was used to define SAF requirements, then operator intervention is required at run-time to provide the appropriate behavior and subsequent training experience.

3. The way behavioral fidelity is defined for a particular application in terms of these categories will strongly influence the technical approach to achieving desired SAF training fidelity. In Figure 4, a technical solution for SAF training fidelity must be capable of generation of behaviors that were not required (and thus not natively available) of native SAFs. At another extreme, illustrated in Figure 5, behavior fidelity includes any realistic behavior. In this figure, a behavioral goal has been identified that would be appropriate to present to the trainee that is realistic and functional but not doctrinal. In this case, technology to achieve SAF training fidelity would focus on selection of behaviors because the realism requirement had already provided the ability to present the targeted behavior. In general, selection is much less technically difficult than generation (Bonasso et al., 1997).

In practice, solutions that attempt to improve SAF training fidelity are likely to require a mix of generation and selection. From a requirements analysis and definition point of view, requiring that SAFs provide some realistic but non-doctrinal behaviors to support selection may result in less development expense and reduce technical risk when greater SAF training fidelity is desired.

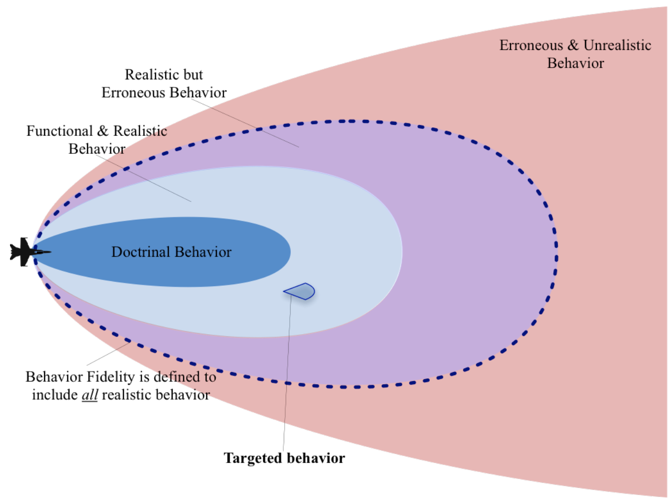

Figure 5. SAF Behavioral Realism Influences Technical Requirements for Improved SAF Training Fidelity.

Representational Requirements

To be able to deliver the required SAF training fidelity, the system will need to have some understanding of the training context and mechanisms to adapt experience. These are representational requirements because the knowledge or ability must be present in the system in some form to improve SAF training fidelity. Options for specific form(s) of the representation(s) can be evaluated against the requirements of a particular system to determine which representations are more apt for that system.

Table 2 summarizes the representational requirements for SAF training fidelity. The table divides the requirements into functional and user-interaction categories. Functional requirements are the ones that are strictly required to enable SAF training fidelity (i.e., targeting of training goals and adaptation based on proficiency with reduced human intervention). User-interaction requirements may not be strictly necessary. However, our experiences prototyping and testing technologies that attempt to improve SAF training fidelity (Schatz et al., 2012; Wray et al., 2009; Wray & Woods, 2013) suggest that providing users (instructors and operators) with the ability to understand, predict, control, and customize an adaptive learning system is critical for successful adoption. These instructor-mediated design requirements (Folsom-Kovarik, Wray, & Hamel, 2013; Wray & Munro, 2012) are consistent with other observations about the importance of instructor control for effective deployment of adaptive learning systems.

SYSTEMS ENGINEERING & SAF TRAINING FIDELITY

Systems engineering concerns how the SAF training fidelity requirements are incorporated into the design and implementation of the overall training system. Figure 6 suggests some of the system design alternatives that will need to be evaluated in systems designed to improve SAF training fidelity. We focus on options for improving SAF training fidelity in existing or typical simulation-based training systems.

The components that support improved SAF training fidelity are filled with gray; components normally present in the training environment are show without fill. Information from the trainee (e.g., interacting in a virtual cockpit) flows to both the simulation engine and a "decision" component that evaluates the current situation in terms of training goals and proficiency. The decision component is likely to interface to the system in the same ways, regardless of the way its decisions are carried out in the training system.

Table 2. Representational Requirements for Improving SAF Training Fidelity.

Requirement	Rationale
Functional requirements	
Training goals	The system requires some representation of the training objectives in order to be able to target goals and re-redirect/repair when progress toward the goal is stopped. Explicit, semantically rich representations of training goals and scenario events can support adaptive training.
Trainee proficiency	Although adaptive tailoring does not necessarily require a learner model (Durlach & Spain, 2012), the system needs some method for associating trainee actions with estimated levels of proficiency. This requirement can be met with traditional proficiency models (Dillenbourg & Self, 1992; Foster & Fletcher, 2003) but perhaps also with active-learning data-mining that reduces the need for a priori models without needing millions of examples (Pan & Yang, 2010; Settles, 2012)
Mechanisms of adaptation	The system requires some algorithm or methodology for generating and/or selecting behavioral options or actions during run-time that would improve overall SAF training fidelity. Many technologies and methods may be apt including planning, rule-based systems, probabilistic models (POMDPs), etc. The choice of a particular technology will be influenced by the anticipated mix of generation and selection (as above).
User interaction requirements	
Functional abstractions	Most current SAF implementations expose some behavioral options to an operator/instructor to enable parameterization for a mission. SAF training fidelity will require additional parameters, which should be exposed to the user. Because users will often not have as much experience and knowledge of instructional design (as compared to the domain), the customization parameters should be formulated at a level of abstraction that users will readily understand. For example, allowing the user to set initial levels of "helpfulness" and "task complexity" is likely to be better than detailed choices for instructional strategies and scaffolding options.
Alternative modalities	Different modes of employment and use will result in different requirements for the mode of operation. Potential modes include enabling/disabling, recommendation (e.g., bringing operator attention to potential need and recommending options), autonomous operation, and mixed-initiative or supervisory control paradigms (Sheridan, 2002) where the system has some pre-specified capability to act as well as limits on the contexts and actions it is allowed to execute.
Explicit rationales	As suggested in the introduction, after-action explanations of SAF decision-making are often required for training. When a system makes choices that are motivated by training goals, explanation of rationales is likely even more important. For example, if a situation that occurred relatively infrequently in the environment was presented to the trainee to support a training goal, helping the trainee and instructor both understand the reasons it was introduced and explicitly discussing its rarity would help the trainee better situate the experience in the training program. Explanations and rationales to instructors are also likely to improve understandability and trust.

Figure 6a illustrates the technically simplest option, which is to bias or over-ride the behavior selection mechanism in the native SAF. This option requires no behavior generation capability other than that provided by the SAFs. Importantly, this option does not assume that the behavior realism requirements are defined broadly, as in Figure 5. It could be used even with the doctrinal definition of behavior fidelity in Figure 4. In this case, however, selection would be limited to options available from doctrine. The main limitation of the Figure 6a approach is that improvements to SAF training fidelity are bounded by the range of native SAF behaviors. Figure 6b suggests an approach that allows extension of doctrinal behaviors by the inclusion of SAF behaviors that are non-doctrinal. These behaviors could be implemented a priori, resulting in a hybrid SAF model (Wray et al., 2005). They could be generated dynamically, although this would be more technically challenging.

Figure 6c illustrates a choice that offers a compromise between these two design options. In this approach, SAFs are directed to perform non-native actions. The SAF training fidelity components act as a director (Magerko, 2007) or command-and-control/operator (Wray, et al., 2005) for the SAFs. This approach requires that the SAFs offer some program-level interfaces for command and control (i.e., not just control at the level of the graphical user interface). This approach can use a pre-existing "behavior library" to deviate from native behaviors or it could generate a control plan. The key difference is that this approach uses controls within the SAF to over-ride native behavior generation rather than direct generation of SAF behavior.

Figure 6c represents the approach we are taking to demonstrate improvements in SAF training fidelity in specific training systems (Wray & Woods, 2013). For example, the Training Executive Agent (TXA, Wray et al., 2015) integrates with a Navy threat system to allow pre-authored interventions ("directives") to over-ride native threat system SAF behavior. In the successive intercept example from above, the TXA is used to define a directive ("reverse-leash") that will ensure that the second group stays within an instructor-defined range boundary (e.g., 70-90 miles) of the first group. When the trailing group nears the edge of the boundary, the directive provides new heading and speed commands to the SAF, ensuring that it remains within the range bound. This behavior ensures that the training goal of the successive intercept can be presented, regardless of how the trainee conducts an intercept against the leading threat.

The hypothesis motivating the Figure 6c approach is that SAF command-and-control functions provide a mechanism for improving SAF training fidelity without requiring substantial engineering impacts on the existing system. This was accomplished with the TXA implementation, which is taking advantage of native SAF capabilities but enabling new behaviors that would have previously required run-time operator intervention. Thus, for systems that already have developed SAFs (with whatever level of realism that was defined originally for those systems), this approach offers a path to improve SAF training fidelity while leveraging prior investments in the SAF and also minimizing new engineering requirements for the SAF.

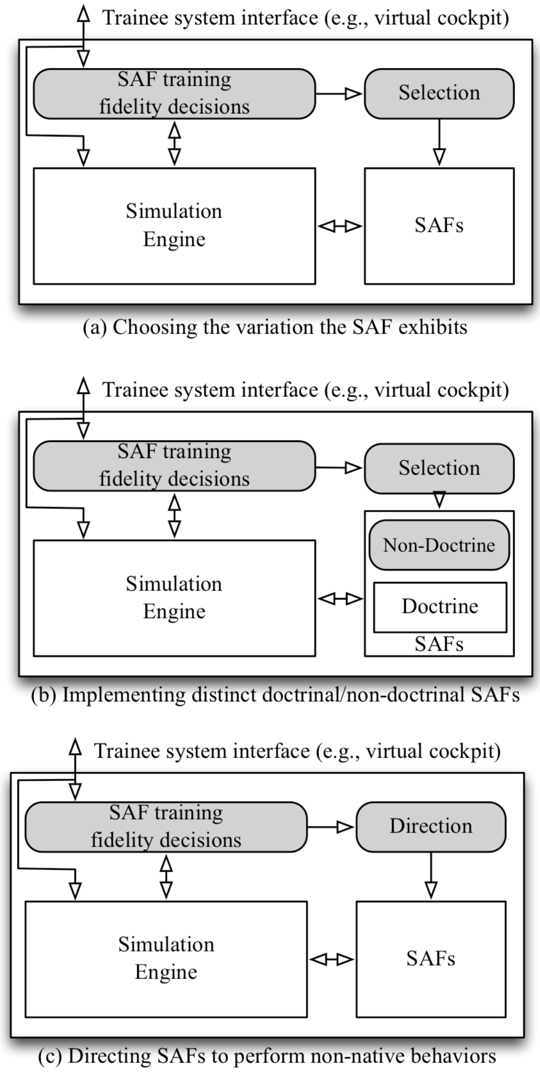

Figure 6. Systems Design Options.

EVALUATING SAF TRAINING FIDELITY

Evaluating SAFs, like the evaluation of any general-purpose tool, is difficult (Campbell & Bolton, 2005). SAF training fidelity requirements do add complexity to the challenges of verification and validation of SAFs generally and also for a particular application. Here we briefly discuss some suggestions for evaluating SAF training fidelity within the context of evaluation of SAFs more generally.

Verification is the process of determining if the implementation of a capability is consistent with the design. For SAF training fidelity, verification requires demonstrating, in a systematic way, that the system responds differently to similar mission situations when the training context differs. We have developed a methodology for verifying such functionality and used it to verify the TXA (Wray, et al., 2015). The approach requires:

- **Simulated trainees**: SAF training fidelity responses depend on trainee actions. Thus, stimulating the system as a trainee might, but in some systematic way, is needed. Simulated trainees, based on subject-matter-expert (SME) analysis of common patterns of trainee behavior and implemented as SAFs in the training simulation, allow automated exploration of the system response space.
- **Explicit representation of training presentation quality**: SMEs provide quantitative criteria that describe goals and trade-offs in particular training scenarios, including variations based on proficiency. These criteria then provide a basis for automating scoring of scenario results in terms of the criteria (Jones et al., 2015).

- **Scenarios of varying complexity**: Multiple scenarios help ensure that the generality of the implementation is assessed. Varying the complexity of the test scenarios helps to identify where automated methods of improving SAF training fidelity can reduce operator workload and/or reduce complexity limitations for more realistic training contexts (Wray, et al., 2015).

Validation is the process of determining if the implementation of a capability meets the requirements. For SAF training fidelity, validation studies would demonstrate that actual training delivers more training experiences that are targeted to training goals and adaptive to trainee proficiency while not increasing the workload on operators and instructors. Validation studies are likely to compare system performance with and without SAF training fidelity functions over a desired operating range (e.g., an existing curriculum of training scenarios). This paper has presented the ultimate goal in terms of *improved SAF training fidelity* because validation is likely to be case in terms of improvement in measures overall a baseline (e.g., the system without SAF training fidelity functions).

One of the hypothesized advantages of the verification design outlined above is that the same methodology can be used for validation, using actual trainees for the simulated ones and we are in the process of designing a study to attempt to validate a particular SAF training fidelity technology following this model. This is a potential next step in evaluation of the TXA. Verification is helpful as well to identify representative candidates in the overall space of scenarios where SAF training fidelity improvements are most needed (Wray, et al., 2015).

CONCLUSIONS

SAFs need to take the training context, not just tactical context, into account when generating behavior. Doing so will contribute to the long-time goal of fully autonomous, intelligent forces for training simulation. This paper introduces issues relevant to *SAF training fidelity in constructive simulation*, describing requirements for SAF behavior generation that go beyond realistic, tactically-appropriate behavior. We introduced rationales, functional requirements, systems design considerations, and evaluation methods for improved SAF training fidelity, focusing especially on achieving this goal in today's SAFs and simulations. For the longer term, identifying SAF training fidelity requirements in the overall system requirements process will potentially result in more capable and less operationally-costly simulation-based training systems as behavioral needs beyond doctrine are recognized and generated by future SAFs.

ACKNOWLEDGEMENTS

This work is supported in part by the Office of Naval Research. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of Defense or Office of Naval Research. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. This work was supported, in part, by the Office of Naval Research project N00014-1-C-0170 Tactical Semi-Automated Forces for Live, Virtual, and Constructive Training (TACSAF). We would also like to thank colleagues who have contributed to the conceptual development and implementation of the work described herein including Ben Bachelor, Ami Bolton, CDR Joseph Cohn, LCDR Jeff Grubb, Rob Falco, Jeremiah Folsom-Kovarik, John Mooney, Charles Newton, and CDR Brent Olde.

REFERENCES

Bonasso, R. P., Firby, R. J., Gat, E., Kortenkamp, D., Miller, D. P., & Slack, M. G. (1997). Experiences with an Architecture for Intelligent, Reactive Agents. *Journal of Experimental and Theoretical Artificial Intelligence, 9*, 237-256.

Bransford, J. D., & Schwartz, D. L. (1999). Rethinking Transfer: A Simple Proposal With Multiple Implications. In A. Iran-Nejad & P. D. Pearson. (Eds.), *Review of Research in Education* (Vol. 24). American Educational Research Association: Washington, DC.

Campbell, G. E., & Bolton, A. (2005). HBR Validation: Integrating Lessons Learned from Multiple Academic Disciplines, Applied Communities and the AMBR Project. In K. Gluck & R. Pew (Eds.), *Modeling Human Behavior with Integrated Cognitive Architectures: Comparison, Evaluation, and Validation*. Matawan, NJ: Lawrence-Erlbaum Associates.

Ceranowicz, A. (1994). Modular Semi-Autonomous Forces. In J. D. Tew, S. Manivannan, D. A. Sadowski & A. F. Seila (Eds.), *Proceedings of the 1994 Winter Simulation Conference*.

Chi, M. T. H., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from Human Tutoring. *Cognitive Science, 25*(4), 471-533.

Dillenbourg, P., & Self, J. (1992). A framework for learner modeling. *Interactive Learning Environments, 2*(2), 111-137.

Durlach, P., & Spain, R. (2012). Framework for Instructional Technology. In V. G. Duffy (Ed.), *Advances in Applied Human Modeling and Simulation*: CRC Press.

Ericsson, K. A., Krampe, R. T., & Tesch-Romer, C. (1993). The Role of Deliberate Practice in the Acquisition of Expert Performance. *Psychological Review, 100*, 3.

Folsom-Kovarik, J. T., Wray, R. E., & Hamel, L. (2013). *Adaptive Assessment in an Instructor-Mediated System*. Paper presented at the Artificial Intelligence in Education (AIED).

Foster, R. E., & Fletcher, J. D. (2003). *Modeling the user for education, training, and performance aiding*. Paper presented at the Proceedings of the NATO Research and Technology Agency. Human Factors and Medicine Panel; 2003 Symposium on Advanced Technologies for Military Training.

Gluck, K., & Pew, R. (Eds.). (2005). *Modeling Human Behavior with Integrated Cognitive Architectures: Comparison, Evaluation, and Validation*. Matawan, NJ: Lawrence-Erlbaum Associates.

Jones, R. M., Bachelor, B., Stacy, W., Colonna-Romano, J., & Wray, R. E. (2015). *Automated Monitoring and Evaluation of Expected Behavior*. Paper presented at the International Conference on Artificial Intelligence.

Magerko, B. (2007). Evaluating Preemptive Story Direction in the Interactive Drama Architecture. *Journal of Game Development, 3*.

Mayer, R. E. (2008). *Learning and Instruction*. Upper Saddle River, New Jersey: Pearson Education, Inc.

Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering, 22*(10), 1345-1359.

Pea, R. D. (2004). The social and technological dimensions of scaffolding and related theoretical concepts for learning, education, and human activity. *The Journal of the Learning Sciences, 13*(3), 423-451.

Pew, R., & Mavor, A. (Eds.). (1998). *Modeling Human and Organizational Behavior: Application to Military Simulations*. Washington, DC: National Academies Press.

Scales, R. H. (2006). The Second Learning Revolution. *Military Review, 86*((Jan-Feb)), 37-44.

Schatz, S., Wray, R., Folsom-Kovarik, J. T., & Nicholson, D. (2012). *Adaptive Perceptual Training in a Virtual Environment*. Paper presented at the Human Factors and Ergonomic Systems (HFES-2012).

Settles, B. (2012). Active Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning, 6*(1).

Shaw, R. L. (1985). *Fighter Combat: Tactics and Manuevering*. Annapolis, MD: Naval Institute Press.

Sheridan, T. B. (2002). *Humans and automation: System design and research issues*: Wiley.

Shute, V. J., & Zapata-Rivera, D. (2012). Adaptive Educational Systems. In P. J. Durlach & A. M. Lesgold (Eds.), *Adaptive Technologies for Training and Education* (pp. 7-27). New York: Cambridge University Press.

Stacy, W., Walwanis, M. M., Wiggins, S., & Bolton, A. (2013). Layered Fidelity: An Approach to Characterizing Training Environments *Proceedings of the 2013 Inter-service/Industry Training, Simulation, and Education Conference (I/ITSEC 2013)*. Orlando: NTSA.

Wallace, S., & Laird, J. E. (2003). *Behavior Bounding: Toward Effective Comparisons of Agents & Humans.* Paper presented at the Eighteenth International Joint Conference on Artificial Intelligence, Acapulco, Mexico.

Wray, R. E., Bachelor, B., Jones, R. M., & Newton, C. (2015). Bracketing human performance to support automation for workload reduction: A case study *Lecture Notes in Computer Science: Proceedings of the Human Computer Interaction International (HCII) Conference*. Los Angeles: Springer-Verlag.

Wray, R. E., Lane, H. C., Stensrud, B., Core, M., Hamel, L., & Forbell, E. (2009). *Pedagogical experience manipulation for cultural learning*. Paper presented at the Workshop on Culturally-Aware Tutoring Systems at the AI in Education Conference

Wray, R. E., & Munro, A. (2012). *Simulation2Instruction: Using Simulation In All Phases of Instruction*. Paper presented at the 2012 Interservice/Industry Training, Simulation, and Education Conference.

Wray, R. E., van Lent, M., Beard, J. T., & Brobst, P. (2005). *The Design Space of Control Options for AIs in Computer Games.* Paper presented at the IJCAI 2005 Workshop on Reasoning, Representation, and Learning in Computer Games, Edinburgh, Scotland.

Wray, R. E., & Woods, A. (2013). A Cognitive Systems Approach to Tailoring Learner Practice. In J. Laird & M. Klenk (Eds.), *Proceedings of the Second Advances in Cognitive Systems Conference*. Baltimore, MD.