A Decision Aid for Optimizing Experimental Design Involving LVC Environments

Sylvain Bruni, Kenyon Riddle, Andres Ortiz, Danielle Dumond, Spencer Lynch
Aptima, Inc.
Woburn, MA
sbruni@aptima.com, kriddle@aptima.com, aortiz@aptima.com, ddumond@aptima.com, slynch@aptima.com

Henry Marshall, Chris Gaughan
U.S. Army Research Laboratory Human Research and Engineering Directorate Simulation and Training Technology Center
Orlando, FL
henry.a.marshall@us.army.mil, chris.gaughan@us.army.mil

Jay Saffold
Research Network, Inc.
Kennesaw, GA
jsaffold@resrchnet.com

ABSTRACT

Increasingly, Modeling and Simulation (M&S) is playing a key part in the decisional process Program Managers (PM) make in the development of new systems, testing, doctrine, and other processes. Unfortunately, the PM must navigate their decisions about leveraging M&S without any supporting aids, making entry and efficient utilization difficult. There is currently no systematic method for assembling environments and designing experiments from multiple M&S perspectives like Live vs. Virtual vs. Constructive simulations to provide decisional data. This process typically requires multiple stakeholders to meet many times in an effort to assemble modeling and simulation-based experiments "that work." As more models, simulators, and scenarios become networked and available to experimenters, a solution is needed to facilitate and accelerate the setup of complex experiments that involve these assets. To meet this need, research was conducted to develop the Live Virtual Constructive & Game - Assisted Experimental Designer tool (LVC&G-AED), an interface and software solution that guides individuals through a ten-step research process, from defining research questions and choosing variables of interest, to developing relevant measures and specifying the environment's software and hardware apparatus. This process is designed to be high-level, capturing the questions of the various professionals involved in simulation development, while being sufficiently rigorous to ensure that specific research questions are addressed. Partially Observable Markov Decision Process algorithms, coupled with an intuitive user interface, allow for interactive exploration of the state space of experimental configurations of simulators, equipment, and other resources available to the user. Through the LVC&G-AED decision-aid, experimenters are provided with recommendations for optimal experimental design configurations. Ultimately, LVC&G-AED translates experimental and simulation requirements into machine-actionable constraints, to facilitate the complex setup of experiments that involve combinations of Live, Virtual, Constructive, and Game M&S environments. This paper focuses on the development lessons learned during this research and the way forward.

ABOUT THE AUTHORS

Sylvain Bruni is a Senior Human Systems Engineer and leads the Cognitive Systems Integration capability lead at Aptima, Inc., where he provides expertise in human-centered design processes: requirements definition, design development and implementation, and system evaluation. His research focuses on human-automation collaborative systems, multimodal user interfaces, human-in-the-loop experimentation, and the statistical analysis of human-centered system data. Mr. Bruni holds a Scientiæ Magister in Aeronautics and Astronautics from MIT and a Diplôme d'Ingénieur from the Ecole Supérieure d'Electricité (Supélec, France). He is a member of the Human Factors and Ergonomics Society, the Institute of Electrical and Electronic Engineers (IEEE) Systems, Man, and Cybernetics Society, the Association for Computing Machinery, and the Department of Defense Human Factors Engineering Technical Advisory Group.

Kenyon Riddle is a Human Factors Scientist and Manager of Orlando Operations at Aptima, Inc. Mr. Riddle provides expertise in the design, evaluation, and implementation of human-centered automation and decision-support tools, as well as in human-system evaluation, experiment design, and statistical analysis. Prior to joining Aptima, Mr. Riddle was a Research Associate at Honeywell Aerospace in Golden Valley, MN, where he worked on a wide range of projects supporting the design and analysis of human-centered systems in the aviation domain. Mr. Riddle holds a M.S. in Human Factors from the University of Illinois at Urbana-Champaign and a B.S. in Psychology from the University of Central Florida. He is a member of the Human Factors and Ergonomics Society, the American Psychological Association, and the Aircraft Owners and Pilots Association.

Henry Marshall is a Science and Technology Manager at the Army Research Laboratory (ARL) Human Research and Engineering Directorate (HRED) Simulation and Training Technology Center (STTC). His assignment experience spans across several agencies including Army, Department of Homeland Security (DHS), and Navy. His 30 years with the Government have been spent assigned to leading edge simulation technology efforts in Modeling and Simulation (M&S) Architecture, law enforcement training, embedded training technology, Semi-Automated Forces (SAF), and simulation software development and acquisition. He received a Bachelor of Science in Engineering degree in Electrical Engineering and a Master of Science degree in Systems Simulation from the University of Central Florida (UCF).

Chris Gaughan is the Chief Engineer for Advanced Simulation and Deputy Technology Program Manager of the Modeling Architecture for Technology, Research and Experimentation (MATREX) program at the United States Army Research Laboratory (ARL), Human Research and Engineering Directorate (HRED), Simulation and Training Technology Center (STTC). He has a diverse portfolio of distributed simulation projects that support the full spectrum of the Department of Defense Acquisition Life Cycle. He received his Master of Science and Bachelor of Science in Electrical Engineering from Drexel University in Philadelphia, PA.

Andres Ortiz is an Aerospace Engineer at Aptima Inc., with a broad background in modeling and analysis of multi-agent autonomous systems. Particular areas of expertise include continuous and discrete event system modeling, real-time task scheduling and assignment, optimization, vehicle trajectory planning, control of dynamical systems and machine learning. Dr. Ortiz's research involves the development of augmentation and support systems that facilitate and enhance interaction and coordination between autonomous/robotic systems and their human operators. Dr. Ortiz received his Ph.D. and M.S. degrees in Aerospace Engineering from the University of Illinois at Urbana-Champaign, and a B.E. in Electronic Engineering from the Pontificia Universidad Javeriana in Cali, Colombia.

Danielle Dumond is an Analytics, Modeling and Simulation Research Engineer at Aptima, Inc. Her background lies in machine learning, decision making, probabilistic modeling and robotics. Dr. Dumond's interests include using mathematical models to predict and control how different agents within a system will behave and interact. Dr. Dumond received a Doctorate of Philosophy (Ph.D.) in Engineering Science with a specialization in Control Engineering from Thayer School of Engineering at Dartmouth College, and a Bachelor of Science in Physics from the College of William and Mary. She is a member of the Society of Women Engineers (SWE).

Spencer Lynch is a Software Engineer at Aptima, Inc. who is skilled in fullstack web development. His experience writing web applications includes writing Representational State Transfer conforming applications using the latest technologies such as WebSockets, Web Graphics Library, Web Real-Time Communication, Data-Drive Documents JavaScript and he has much experience writing applications adhering to responsive design principles. He holds a B.S. in Computer Science and has three years of experience writing web applications and simulation software.

James Saffold is the President, CEO, and Chief Scientist for the Research Network Inc. (RNI). He has over thirty years of experience as an engineer in both the military and industry. He holds a Bachelor of Science in Electrical Engineering degree from Auburn University (1983). Mr. Saffold has performed research in RF Tags, Ultra Wide Band (UWB) radar, Virtual Reality, Digital databases, Soldier Tracking Systems, Millimeter wavelength (MMW) radar, multimode (MMW and optical) sensor fusion, fire-control radar, electronic warfare, survivability, signal processing, and strategic defense architecture. Mr. Saffold's current interests include embedded After Action Review (AAR) systems, game-based training, combat identification, virtual scenario simulation, sensor fusion, signal processing, electromagnetic propagation, and phenomenology. Mr. Saffold lectures annually at the Georgia Institute of Technology on topics related to remote sensing, propagation, clutter, smart munitions, and signal processing.

A Decision Aid for Optimizing Experimental Design Involving LVC Environments

Sylvain Bruni, Kenyon Riddle, Andres Ortiz, Danielle Dumond, Spencer Lynch
Aptima, Inc.
Woburn, MA
sbruni@aptima.com, kriddle@aptima.com, aortiz@aptima.com, ddumond@aptima.com, slynch@aptima.com

Henry Marshall, Chris Gaughan
U.S. Army Research Laboratory Human Research and Engineering Directorate Simulation and Training Technology Center
Orlando, FL
henry.a.marshall@us.army.mil, chris.gaughan@us.army.mil

Jay Saffold
Research Network, Inc.
Kennesaw, GA
jsaffold@resrchnet.com

INTRODUCTION: THE EASE CONCEPT OF OPERATIONS

Increasingly, Modeling and Simulation (M&S) plays a key part in the decisional process Program Managers (PMs) make in the development of new systems, testing, doctrine, and other processes. Unfortunately, the PM must navigate a complex decision space to leverage M&S, without any dedicated supporting aids. The sheer amount of information available and the vast number of considerations that must be taken into account make efficient utilization of this research space difficult. There is currently no systematic method for assembling environments and designing experiments from multiple M&S perspectives such as Live vs. Virtual vs. Constructive simulations to provide decisional data. This process typically requires multiple stakeholders to meet many times to assemble M&S-based experiments that sufficiently answer research questions. As more models, simulators, and scenarios become networked and available to experimenters, a solution to this complex decisional process is needed to facilitate and accelerate the setup of experiments that involve these assets.

From the introduction of new technology to dynamic enemy tactics, the ability to test and train new technologies, operating procedures, and organizational structures in a quick and cost-effective manner is becoming increasingly important. Live, Virtual, Constructive, and Game (LVC&G) simulations allow for efficient and repeatable experimentation; however, many individuals are involved with the creation of LVC&G simulator hardware, software, and scenarios. An Army analyst understands the phenomena to be studied. An operations researcher or systems analyst understands research methods, experimentation, and measurement. A simulation engineer knows the intricacies of the LVC&G scenarios and general capabilities specific to the LVC&G platform. In order to create an appropriate scenario, all three of these individuals must coordinate. The difficulty of coordinating across people and professions can slow development of tests – producing tests that are less decisive and efficient than desired, while adding additional financial and manpower costs (McDonnell et al., 2011).

To bridge this gap, the United States (U.S.) Army's Simulation and Training Technology Center (STTC) has undertaken the conception of an Executable Architecture Systems Engineering (EASE) as a unifying platform that connects stakeholders with all the M&S equipment at their disposal (Marshall, 2011). The fundamental purpose of EASE is to ensure interoperability and connectivity between the users and their tools, in a manner that simplifies access and implementation of experimental or training configurations (Figure 1). Beyond program managers, stakeholders may include M&S users (e.g., experimenters or trainers), system engineers, developers, and other subject matter experts. M&S equipment may refer to various apparatus available on-site or remotely, including models, simulators, scenarios, hardware, software, and data repositories.

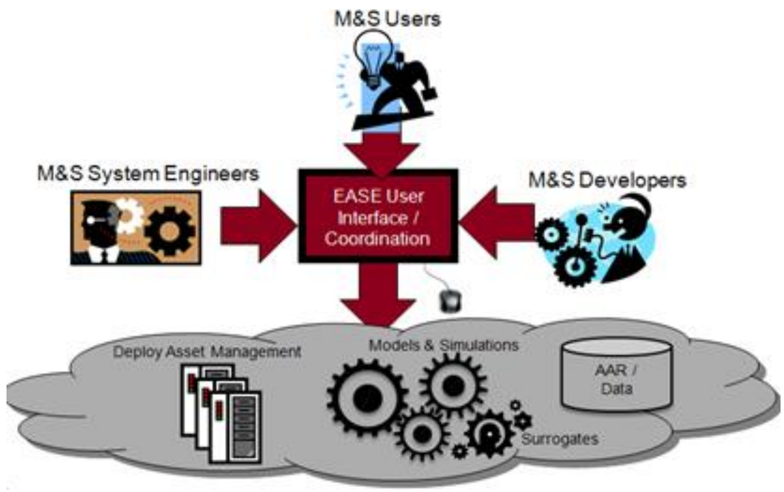

Figure 1. STTC's EASE Concept of Operations

INTRODUCING AUTOMATED RECOMMENDATIONS IN THE DECISION-MAKING PROCESS

The development of a highly automated solution can streamline the selection and setup of simulation environments by guiding users through the experimentation process, while still leveraging much of the knowledge of the stakeholders by utilizing an intelligent solution. The experiment designs developed with such technology would be well-structured and complete, so design teams can coordinate more productively over these products. To test this approach, we developed the LVC&G Assisted Experimental Designer (AED), a machine intelligent and scientifically-informed guided research aid to support a variety of networked LVC&G simulations.

As illustrated in Figure 2, the foundation of the AED decision-support system is a computer-based automated interview process that enables users to input information relevant to their experimental objectives. This interview is guided by a ten-step research questionnaire, structured as a systematic way of designing an experiment, from the identification of research goals and statistical design to the definition of independent and dependent variables. An intelligent algorithm then analyzes the input provided by the experimenter, in the form of a Partially-Observable Markov Decision-Process (POMDP) model, which explores the domain space of feasible experimental configurations that would satisfy the user's objectives. The POMDP model dynamically queries a Knowledge DataBase (KDB) containing information relevant to existing M&S or live assets, platforms or scenarios, as well as records of previously-conducted experiments. With these data, the model infers feasible experimental configurations that optimize the design for cost (i.e., minimizing the cost of running the experiment) and for quality (i.e., maximizing the quality of the expected experimental results). A rank-ordered list of experimental configurations is returned to the experimenter through the AED interface, allowing the user to select the configuration they want to run. Ultimately, AED will be able to automatically push an experimental design file based on the selected configuration directly to the M&S or live platforms connected with EASE.

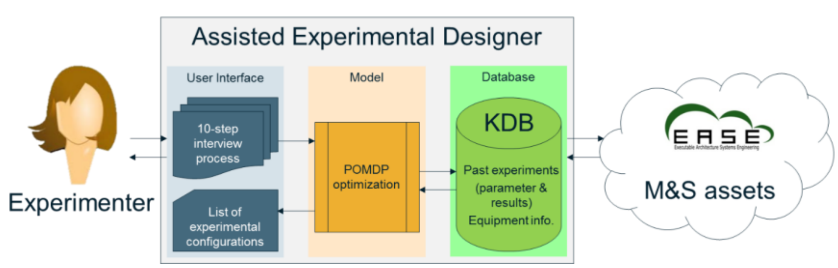

Figure 2. System Architecture for the Assisted Experimental Designer

In order to maximize the benefit of employing automated reasoning, we designed the AED tool such that it could operate in a closed loop, constantly refining its recommendations for experimental configurations based on the incremental new knowledge gained after each experiment.

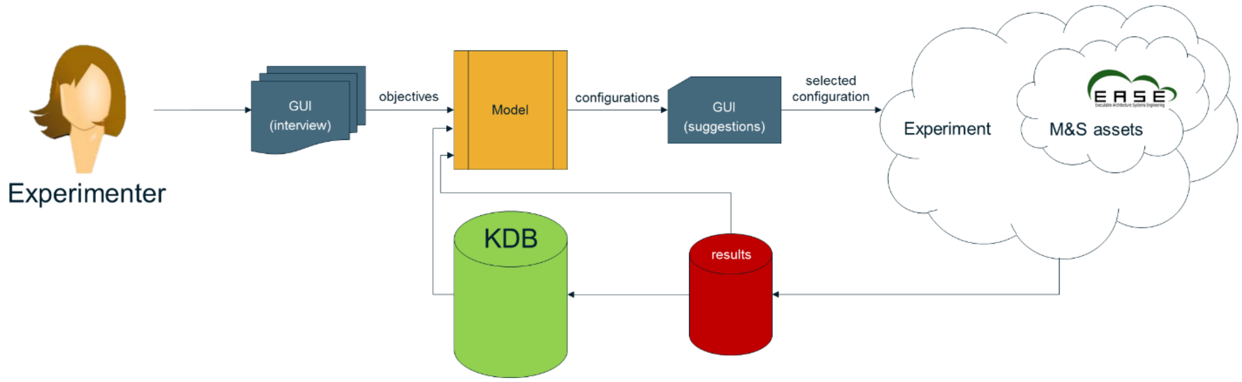

Figure 3. Process-oriented View of AED

The process diagram of Figure 3 illustrates the closed loop permitted by the usage of POMDP algorithms at the core of the AED decision-aid. After an experimental configuration ("experiment") is selected and implemented, AED pushes the resulting data ("results") to the KDB for future use. These results are also directly analyzed by the model, which compares them to past results from the KDB and to the initial objectives input by the experimenter. The model is then capable of updating its recommendations for experimental configurations so that the next experiment may yield better results.

THE GRAPHICAL USER INTERFACE

To facilitate the input of experimental information by a user, an intuitive Graphical User Interface (GUI) was built. It leverages a ten-step experimental design process that queries the user for key information related to the experiment they wish to conduct (Figure 4). The ten-step experimental design process was built from previous versions developed for the Office of Naval Research and for the Air Force Research Laboratory, in research projects related to evaluating operator supervisory control performance in multi-asset resource allocation under uncertainty (Cummings and Bruni, 2010) and to testing various experimental designs for target detection tasks by unmanned systems operators (Swanson et al., 2012). This process was built to ensure that no key experimental parameter is left out during the design and execution of the experiment, and that the design of the experiment is repeatable in a manner that accounts for statistical validity.

Figure 4. Screenshot of the Prototype Graphical User Interface for AED Interview Process

The first seven steps, listed in a quick-access menu in the left panel and available as a form in the main area of the GUI, include: (1) **objectives**, (2) **research questions**, (3) **constraints**, (4) **environment**, (5) **independent variables**, (6) **dependent variables**, and (7) **statistical design**. Once the experimenter has filled out the first seven steps, the model generates an experimental configuration based on user inputs and queries to the knowledge database. The experimental configurations generated by the model are pushed back to the user in (8) **configuration**, a rank-ordered list of experimental configurations generated by the AED model for the experimenter to select from and modify as they see fit (e.g., "*One Semi-Automated Forces (OneSAF) with Scenario 4 and the Vision Toolkit*"). The process ends with two additional steps: (9) **protocol**, where the system outputs an automatically-generated experimental protocol document that lists information pertaining to the experiment with the selected configuration and (10) **Internal Review Board (IRB)**, where the system generates documents pertaining to the submission of an application for IRB approval whenever the research involves human participants.

Figure 5 depicts the GUI for the results input form. While the vision for an integrated version of AED with EASE calls for a direct feed of experimental data and results to be connected to the AED model and KDB, our first implementation of AED requires a manual upload of a comma-delimited file that contains the results of the experiment performed. This current version also includes a section where the experimenter may provide subjective feedback regarding the experiment. An early stakeholder-review of possible dimensions for collection of such feedback has identified a set of costing factors as a primary candidate. AED currently asks the user to rate the costs of personnel, of running the experiment, of acquisition or access to LVC&G assets, and of licensing on a 7-point scale (from "extremely low" to "extremely high"). This information is subsequently stored in the KDB and used by the model to refine its estimates of the cost of feasible experimental configurations. The design of the GUI followed the principles of ecological interface design (Vicente, 2002) and decision-oriented design (Metersky, 1993).

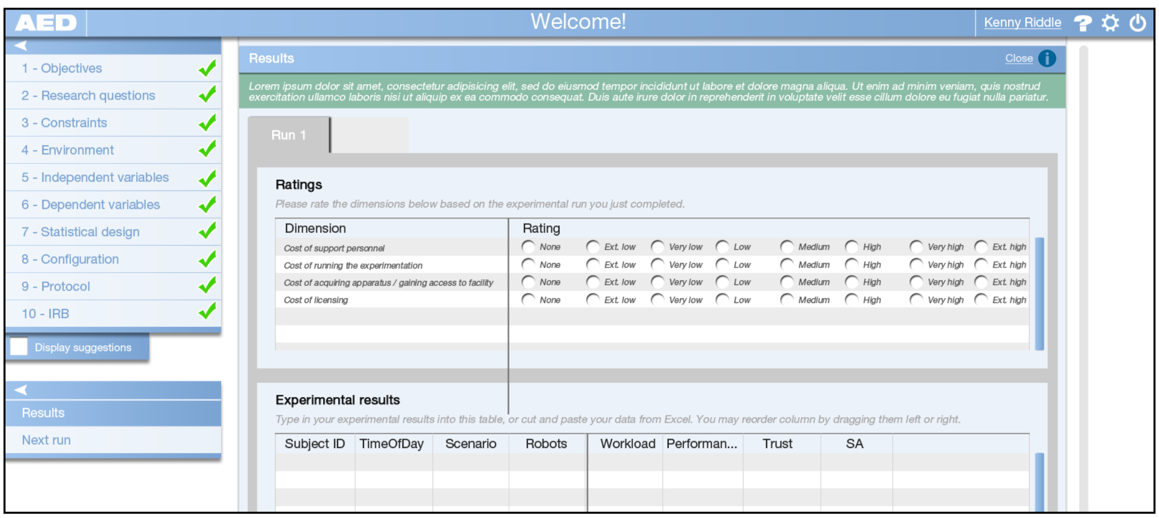

Figure 5. Screenshot of the Prototype Graphical User Interface for AED Results Input Form

THE MODEL

Background

Conducting an experiment is an iterative process in which the experimenter designs the experiment, collects data, and analyzes the results to determine if there is enough evidence to answer the proposed research questions. After each iteration, a decision must be made whether or not to continue based on the quality of the results collected so far and constraints such as time and cost. If the experimenter decides to continue the process, they must modify the current experimental configuration in order to collect data that will ultimately lead to an improvement in the quality of the results. The decision on what to change in the experimental configuration could involve simple adjustments such as increasing the number of trials or randomizing the run order, to more elaborate options such as utilizing a different apparatus.

Given the nature of the experimental process, the AED's decision support system is built by modeling this process as a Partially-Observable Markov Decision Process (POMDP). POMDP models have been widely used to represent and optimize sequential decision-making problems under uncertainty (Smallwood & Sondik, 1973, Puterman, 1994). These types of problems are commonplace in many real-world domains such as robot navigation (Pineau et al., 2007), assistive systems (Hoey et al., 2010), medical diagnosis and treatment (Hauskrecht and Fraser, 2008), inventory management (Treharne and Sox, 2002), and more recently in the domain of automated training and learning systems (Carlin et al., 2013 and Andrews et al., 2013).

The POMDP model provides a powerful tool to capture the uncertainty in determining experimental progress and ultimately provide decision-making support in the form of actionable experimental configurations to novice experimenters. The value of the POMDP over many other modeling techniques is that it allows us to develop a quantitative model of a system, even when certain states of the system are not observable (Puterman, 1994). Traditional Markov processes serve as excellent quantitative models of a system when the transitions of the system from one state to another contain some uncertainty, but these states must be known. When certain characteristics of a system are not directly observable, a POMDP model can utilize those characteristics that *are* observable to develop a probabilistic representation of the system's true state (Smallwood & Sondik, 1973). In the research domain addressed by AED, a system is comprised of every linked simulator apparatus, constructive entities being employed, variables being manipulated and measured, and any humans that may be in the loop. Some of the characteristics that may not be directly observable include the validity of certain constructive entities, the direct relationship between a simulator's fidelity and human performance, and how the modification of variables in one simulator apparatus affects the representation of other variables in a linked apparatus, among countless others. By receiving feedback on results obtained from certain experimental configurations, the POMDP model constantly improves its probabilistic representations of a system's true state.

The POMDP Model Representation

Formally, a POMDP model is defined by a tuple $< S, A, P, \Omega, O, R >$. The set $\boldsymbol{S}$ of all possible states s is referred to as the state space. The set $\boldsymbol{A}$ defines a set of actions while $\boldsymbol{P}$ defines the state transition model. $\boldsymbol{P}$ is a probability table composed of elements $P(s'|s,a)$, which define the probability that the model will transition from state s at time t, into state s' at time $t+1$ given that action a was taken. The set Ω represents a set of possible observations and the observation function $\boldsymbol{O}$ provides the probability that a particular observation will occur given a state transition (i.e., O(o|s',a,s')). Lastly, the reward function $\boldsymbol{R}$ is used to drive the optimization problem, that is, the goal of the POMDP solver is to create a policy to maximize this reward.

The POMDP models the experimental process by tracking experiment state. For this particular application, the state is an abstract representation of the progress towards meeting the experimental goals. More precisely, the state is a representation of the quality of the experimental results that the experimenter has obtained throughout the experimental process. In the current implementation, the state space consists of a vector s which contains a quality value for each one of the independent and dependent variables that are being used to conduct the experiment.

$$s = (q_1, q_2, \dots\dots, q_n) \qquad (1)$$

The true state however, is not directly observable (i.e., measurable) but can be estimated through a set of observations which are related to it. As such, the model combines a set of measurable parameters (i.e., observations) to infer the experimental state based on two categories: (1) the quality of the instrument which is estimated through measures of reliability and fidelity of a particular apparatus for each particular variable; and (2) the quality of the results, which is estimated based on metrics such as mean stabilization, variance and bound checks, outlier reduction, confidence intervals checks for each variable across a number of trials. The observations are calculated after experimental results are available and stored in a vector *o*.

$$o = (o_1, o_2, \dots\dots, o_n) \qquad (2)$$

Where, each observation is a weighted sum of the values for each category.

$$o_i = 0.3 \times Quality\ of\ Instrument + 0.7 \times Quality\ of\ Results \qquad (3)$$

Since the true state is not observable, the POMDP model keeps track of an abstracted version of the state commonly referred to as the belief state b.

$$b = (\tilde{q}_1, \tilde{q}_2, \dots \dots, \tilde{q}_n) \qquad (4)$$

The belief state is a $n \times k$ matrix which contains n vectors $\tilde{q}$, each containing a probability distribution over the k possible values of each of the n states.

$$\tilde{q}_i = (p_1, p_2, \dots \dots, p_k) \text{ and } \textstyle\sum_{i=1}^{k} p_i = 1. \qquad (5)$$

These probability distributions are updated after every observation o and provide a representation of the most likely quality value for each of the state variables. Lastly, the set A, is defined by the courses of action that the experimenter can take after each one of the experimental runs has been completed. These actions are represented in the model as a set of experimental configurations that the experimenter can run. These experimental configurations are obtained using data from historical experiments within the KDB.

$$A = (a_1, a_2, \dots \dots, a_m) \qquad (6)$$

Ultimately, the goal of the model is to recommend to the experimenter the next experiment configuration that should be run to improve upon its current results. In POMDP terms, this translates into selecting an action, given the most current belief state and observation such that the reward function is maximized. This is done by mapping observations to actions using a policy, developed to maximize the reward.

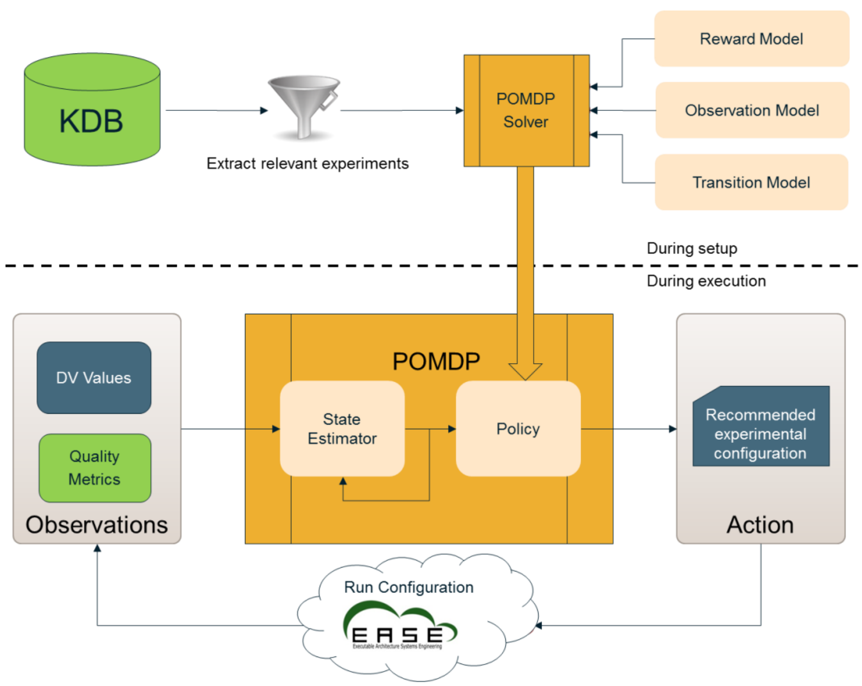

Figure 6. Detailed View of the POMDP Model in AED

In order to achieve this goal, the AED model works in two stages (Figure 6). In the setup stage, a POMDP solver is used to generate a POMDP policy for the execution stage to select the next action. In order to generate the policy the state transition model ***P*** and the observation model ***O*** must be defined for a particular experiment. This is done by characterizing the effect of each possible action (i.e., available and relevant experiment configuration from KDB) on each one of the components of the state (i.e., the IVs and DVs) via two parameters: (1) the *applicability* which refers to how likely the experiment is to increase the quality of a state variable and (2) the *difficulty* which captures how costly it is to achieve that increase in quality. For example, it is possible that a particular experiment produces high quality results for a given DV, in which case the applicability level for this DV would be high. However, it can also be the case that the cost of using the apparatus for this experiment is also high, thereby making the difficulty value high and thus not necessarily the best option. By capturing these parameters, the POMDP solver can create a policy that will select actions with maximum applicability and minimum difficulty across all variables. Note that the larger the action space, the more complex the policy generation stage becomes. In order to keep this complexity manageable the action space is pruned prior to policy generation by filtering past experimental configurations to select only those which are similar to the initial configuration that was selected by the experimenter. The similarity criterion as well as the threshold for selection, are parameters which may be specified by the experimenter.

Once the POMDP policy is generated, the model enters the execution stage. In this stage, new recommendations are suggested to the experimenter every time the experimenter can provide a complete set of results from an experimental run. These results include the measured values for all of the dependent variables of interest for a specific set of trials as well as the values of the independent variables used for each trial. Given these values, observations are generated (based on the aforementioned quality metrics) for the POMDP model to update its state. The new estimated state (i.e., the belief state) is then used by the policy to generate a new recommendation. This iterative process may continue until the experimenter is satisfied with the outcome of the results.

Building the POMDP Model

The AED model was built within a custom developed C# (C sharp) application. The model consists of a POMDP library wrapped inside the decision support module referred to as the AED Engine. We envision the AED to be a service which can be used by several users simultaneously while running distinct sets of experiments with a variety of custom settings. In order to enable such deployment an application server was developed which facilitates the interaction between the multiple modules within AED. Within this architecture, the model operates as follows. An instance of the AED Engine module is created for each user that is carrying out an experimental process. Each AED Engine instance has its associated POMDP module and is managed by the AED Object Manager (AOM). The AOM loads pre-existing experiments from KDB and makes them available in memory to all instances of the AED Engine. Similarly, after every experiment is run, the information from that experiment is stored in the KDB through the AOM. Lastly, the AOM facilitates data passing between the AED Engine and the GUI in order to provide the model with the necessary configuration parameters from the user and the experimental process and return the generated recommendations.

The functionality of the various modules has been initially tested to ensure that all of the pieces deliver the specified functionality and communicate appropriately. Features like loading and saving KDB data, creating experiment configurations from user defined parameters, generating POMDP policies from a set of experimental configurations and generating POMDP recommendations from observations have all been successfully tested by populating the KDB with scripted apparatus and experiment data. In terms of the validity of the recommendation engine itself (i.e., the POMDP model) it is a challenging task to benchmark its performance given the nature of the decision making problem. That is, the potentially large decision space which can lead to a series of recommendation sequences that may be equally sensible. Despite the limited availability of data and the challenge of generating testing scenarios, a few scripted test cases have been developed in order to help determine the validity of the recommendations. It is anticipated that the fidelity of the testing and evaluation process will increase as the KDB gets populated with real historical data. In addition, expertise from experimental designers can be leveraged to determine qualitatively the efficacy and sensibility of the recommendations in a systematic manner.

THE KNOWLEDGE DATABASE

The purpose of the KDB is to act as an information repository available to the model for query when building and evaluating experimental configurations. The KDB can fulfill four queries from the model (see Bruni et al., 2014 for more details):

1. **Request for experimental information**: the KDB houses known information about previously conducted experiments. This information is captured in the KDB in a systematic structure using an Extensible Markup Language (XML) schema (TestExperiment.xml).

2. **Request for apparatus information**: the KDB houses known information about apparatus available to the user for experimentation. This information is captured in the KDB in a systematic structure using an XML schema (TestApparatus.xml).
3. **Request for cost estimate**: the KDB can compute the estimated experimental cost of an experimental configuration designed by the model.
4. **Request for quality estimate**: the KDB can compute the estimated apparatus and experimental quality of an experimental configuration designed by the model.

Using structured interviews to populate the database

In order to fulfill these requests, the KDB leverages a body of data collected directly from current experiments using the AED tool, but also through structured interviews with stakeholders. A survey questionnaire was built to elicit information from expert interviewees, in a manner that mimics the schema of the KDB, but in "human-friendly" questions with example answers. An example question is described below.

> *Example query:* Can the apparatus partition areas (cordon) for experiments or training exercises automatically and if so how is this done and are the partitions modifiable at runtime?
>
> *Example response:* Yes, by setting up route systems in scenarios and/or implementation of entity triggers which can be changed at runtime by the real-time editor (Quality: medium, Confidence: high)

The information is then recorded in XML format in the KDB. Note that the responder's confidence in their answer is always requested during the structured interviews, in order for the model to make an estimate of the associated quality of the characteristic, and the utility of the response in choosing this apparatus for an experimental configuration.

Current status of the KDB

The KDB is currently migrating from an XML-based schema to a Structured Query Language (SQL) database. Initially, an XML configuration file was created for any new database entry (including experiments and apparatus). This allowed for flexibility in the database structure as initial data was collected from stakeholders. As the database has matured and is populated by more entries, the XML schema has been refined to include more information that is beneficial for the POMDP model, while excluding certain information that was found to not be beneficial. Now that the schema has been refined, a proper SQL database will contain the information formerly included in the XML configuration files. As the database grows, this will allow for more computational efficiency, which will allow users to obtain recommendations from the POMDP model rapidly.

DEPLOYMENT AND TESTING AT THE UNITED STATES MILITARY ACADEMY

An initial deployment of the AED system is ongoing at The United States Military Academy (USMA). EASE has been operable at USMA for some time, making this an ideal facility for initial integration efforts between the AED system and EASE. As part of this integration process, knowledge elicitation sessions with various stakeholders and potential end users at USMA have been ongoing to maximize the benefits obtained of this AED-EASE integration. Additionally, insights from potential end users during these knowledge elicitation sessions have allowed the user interface to be modified to better support these users.

This initial integration with EASE could help set a precedent for the future integration of other systems with EASE and could inform the development of an EASE Application Programming Interface (API) in the future. The lessons learned during this integration can be applied at other EASE-equipped facilities, to ease future integration efforts. Additionally, the installation of the AED system at USMA will help gain access to additional stakeholders. Many Army facilities work with USMA to conduct M&S experiments, so this added visibility for the AED system will be very beneficial. Access to more stakeholders will ultimately lead to a larger, more complete KDB, which will only increase the effectiveness of the underlying POMDP model and extend the benefits of the AED system.

CONCLUSION

The purpose of the EASE platform developed by STTC is to ensure interoperability and connectivity between M&S users and their tools, in a manner that simplifies access and implementation of experimental or training configurations. With the Assisted Experimental Designer as a front-end component of EASE, not only are M&S users afforded an intuitive interface to employ EASE and build their experiments, but they also benefit from advanced automation that permits them to make the best use possible of those LVC&G assets connected to EASE. In other words, AED increases accessibility to numerous, distributed modeling and simulation platforms while ensuring more reliable and more efficient M&S or live experiments are designed.

In its initial build, AED relies on POMDP algorithms that explore the space of possible experiments (as captured in the knowledge database) to reconcile feasible configurations with those objectives requested by users, and produce recommendations that maximize the anticipated quality of the experiment while minimizing its cost. Used in series of experiments, AED ingests the resulting data from experiment "n-1" to refine its recommendations for experiment "n," so as to yield significant improvements in experimental outcomes.

Ultimately, EASE combined with the AED decision aid has the potential to augment considerably the capabilities of an experimenter using LVC&G assets. It is envisioned that, through networks of modeling and simulation assets, an experimenter will be allowed to build experiments that make use of those previously inaccessible, remote assets: the AED tool helps to figure out what piece of equipment is the best suited to fulfill the desired experimental objectives, therefore enabling experimenters in one location to drive optimized experiments (defined by AED) in a distributed fashion, using M&S or live apparatus in various remote locations.

ACKNOWLEDGEMENTS

This research was conducted under the Small Business Innovative Research Program from the Army's Simulation and Training Technology Center (STTC). The authors thank Christopher Gaughan of STTC for his support and guidance in this research. The views and conclusions presented in this paper are those of the authors and do not represent an official opinion, expressed or implied, of STTC, the U.S. Army, the Department of Defense or the United States Government.

REFERENCES

Andrews, D., Freeman, J., Andre, T., Feeney, J., Carlin, A., Fidopiastis, C., & Fitzgerald, P. (2013). Training organizational supervisors to detect and prevent cyber insider threats: two approaches. ICST Transactions of Security and Safety.

Bruni, S., Riddle, K., Ortiz, A., Dumond, D., & Saffold, J. (2014). The Assisted Experimental Designer: A Decision Support System to Optimize Modeling and Simulation Experiment Design. *MODSIM World Conference,* April 2014, Hampton, VA.

Carlin, A., Dumond, D., Dean, C., & Freeman, J. (2013). Higher Automated Learning. *Artificial Intelligence in Education (AIED),* July 2013, Memphis, TN.

Cummings, M.L. & Bruni, S. (2010). Human-Automation Collaboration in Complex Multivariate Resource Allocation Decision Support Systems. *International Journal of Intelligent Decision Technologies*, (2010) Vol. 4(2), p. 101-114.

Hauskrecht, M, & Fraser, H. (2008). Planning medical therapy using partially observable Markov decision processes. *Proceedings of the 9th International Workshop on Principles of Diagnosis (DX-98)*, Cape Cod, MA, 182-189, June 1998.

Hoey, J., Poupart, P., von Bertoldi, A., Craig, T., Boutilier, C., & Mihailidis, A. (2010). Automated handwashing assistance for persons with dementia using video and a partially observable markov decision process. *Computer Vision and Image Understanding, 114(5)*, 503–519.

Northeastern University (2014). *How to Calculate Your GPA/QPA and Earned Hours*. Retrieved March 14th, 2014 from http://www.northeastern.edu/registrar/gradecalc.html.

Marshall, H. (2011). Executable Architecture Systems Engineering (EASE), A New Vision for M&S. Proceedings of the INFORMS 2011 Annual Conference, Charlotte, NC.

McDonnell, J.S, Gallant, S., Gaughan, C., & McGlynn L. (2012). Executable Architecture Systems Engineering (EASE). Proceedings of the Systems Engineering Conference (SEDC 2012), Washington, DC.

Metersky, M.L (1993). A decision-oriented approach to system design and development. *IEEE Transactions on Systems, Man and Cybernetics, 23:4*.

Pineau, J., & Gordon, G. J. (2007). POMDP planning for robust robot control. In Robotics Research (pp. 69-82). Springer Berlin Heidelberg.

Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.

Smallwood, R. D., & Sondik, E. J. (1973). The optimal control of partially observable Markov processes over a finite horizon. *Operations Research, 21:5*, 1071-1088.

Swanson, L., Jones, E., Riordan, B., Bruni, S., Schurr, N., Sullivan, S., & Lansey, J. (2012). Exploring human error in an RPA target detection task. *Proceedings of the 56th Annual Meeting of the Human Factors and Ergonomics Society (HFES 2012)*, Boston, Massachusetts.

Treharne, J. T., & Sox, C. R. (2002). Adaptive Inventory Control for Non-Stationary Demand and Partial Information. Management Science, 48(5), 607-624.

Vicente, K. J. (2002). Ecological Interface Design: Progress and challenges. *Human Factors, 44*, 62-78.