Behavioral Recognition and Prediction of an Operator Supervising Multiple Heterogeneous Unmanned Vehicles

Yves Boussemart[1], M. L. Cummings[1]

[1] Humans and Automation Laboratory, Massachusetts Institute of Technology
77 Massachusetts Avenue, 33-407, 02139, Cambridge MA, USA
{yves, missyc}@mit.edu

Abstract. The ability to recognize patterns of operator behavior that could lead to poor outcomes is critical to monitoring the overall performance of the human-unmanned system team. We propose a method that relies on Bayesian machine learning in order to automatically derive a set of states that describe the behavior of an operator. More specifically, we use the Hidden Markov Model (HMM) formalism to infer higher cognitive states from observable operator interaction with a computer interface. This allows the categorization of a pattern of action over a probability distribution of possible operator states which can then be correlated with a range of mission outcomes. Moreover, the HMM provides the means to project in the future and therefore aid in prediction, in probabilistic terms, of the future actions of an operator. In this paper, we present the methodology used to derive the operator model and show initial results based on experimental data.

Keywords: Human Supervisory Control, Hidden Markov Models, Real-Time Behavioral Profiling

1 Introduction

While Unmanned Vehicle (UV) operations currently require multiple operators to control a single platform, future operational paradigms call for a inversion of this ratio by having a single operator controlling multiple unmanned systems, which need not all be of the same type [1, 2]. This, however, increases the number of information sources, the volume of information to be processed and operational tempo, thereby leading to an increased cognitive load on the supervisory human controller [3]. This increased mental workload, in turn, influences the behavioral response of the controller to external events.

The behavior of UV operators is generally guided by Standard Operating Procedures (SOPs). Controllers are trained to perform a specific set of actions in a given situation. During nominal situations, supervisory behavior typically consists of monitoring and/or minimal interaction with the supervised system, such as information request or parameter adjustments. In abnormal situations, however, the controller usually needs to take active measures so as to ensure the safe behavior of the system and/or respond to contingency events that threaten mission success.

While external pressure affects individuals differently, previous research highlights that excessive levels of stress impact the human behavior negatively, and sometimes results in tragic incidents and loss of lives [4]. Recognizing the onset of abnormal behaviors, defined by deviation from the expected behavior, allows for detection and prediction of the occurrence of potential critical events.

This paper describes an effort to model the behavior of an individual controller of multiple heterogeneous unmanned vehicles. The model is based on the Hidden Markov Models formalism and is able to recognize the current behavior of the controller and to forecast the next actions of the controller. More specifically, the model relies on user interface event as an input stream to predict future behaviors.

2 Models of Individual Behavior

Models of human behavior have been of great interest to applied psychologists, artificial intelligence researchers and human factors engineers, all sharing the similar goal of abstracting human behavior into a parsimonious formalism. The vast majority of the work so far has focused on verifying formal theoretical architectures through experimentation (top down).

Applied psychologists traditionally focus on the theoretical aspects of human decision making, notably in terms of preference and bounded rationality [5], multi-attribute theory [6], decision making under uncertainty [7]. Such theoretical notions are usually verified through synthetic experimentation where subjects have to make a choice between multiple options.

Other researchers have attempted to derive computational models of human cognitive processes in order to emulate such processes in computer programs. The commonly used benchmark for such progress is the Turing test, which consists of creating a machine that would exhibit a behavior indistinguishable from that of a human. While the success on the Turing test remains elusive, some researchers restricted their endeavor to expert behavior. Expert systems thus were designed to encapsulate a set of rule abstracted from the knowledge of a human expert [8]. The formalism of agent based modeling has also been used in order to model personality and culture through the use of Goal, Standard and Preference (GSP) trees [9]. Other researchers recognize the limitations of strict rule-based processes and now try to mimic the human common sense in machines [10]. In most of the previous work, the main goal of all these work is to simplify human behavior so that it could be replicated by a machine.

In a more applied setting, some researchers attempt to model human behavior in order to improve performance. For example, theoretical frameworks predicting the consequence of high workload and time pressure are of great interest [11, 12] as they permit an increase of human productivity. GOMS [13] and ACT-R [14] are two cognitive frameworks that try to model the human behavior and that are also used to analyze and improve worker productivity. However, GOMS is limited because it assumes that all users are deterministic and follow the same human processor model. The use of ACT-R is also limited because of sophisticated cognitive task modeling required to fit the framework.

The approach proposed used in this paper differs from all of the above in that it is inherently context and data driven, and does not rely on an a priori theoretical framework of human behavior. The premises of this presented work are akin those of predictive analytics, a data-mining subfield that relies on extremely large corpuses of data to extract patterns in human behavior. While predictive analytics has been used mostly by financial institutions for credit scoring and in other industries such as insurance, telecommunication, retail or travel, we make use of similar data-mining techniques to extract regularities in the interactions of UV supervisory controllers with computer interfaces.

We assume that UV controllers are, because of their training, experts at supervising their systems. Previous work suggests that experts tend to rely on recognition-primed decision-making [15] and therefore tend to exhibit stronger behavioral patterns than novices [16]. We therefore assume that UV controllers similarly exhibit behavioral patterns. The obvious limitation of the data-driven approach is a lack of generalizability beyond supervisory control domains with significant human-computer interaction. However, there are a significant number of applications that would benefit beyond human-unmanned vehicles operations such as air traffic control and first responder settings. Another limitation is that the models are data-driven and thus inherently descriptive: they do not prescribe what the optimal behavior should be without additional a priori performance knowledge. However, the models we propose can prescribe what the "typical" response of a person should be, given prior data about such typical behaviors.

3 Hidden Markov Models

Hidden Markov Models (HMM) were formally defined by the seminal paper by Rabiner et al. [17], and essentially consist of doubly stochastic Markov chains. A HMM has a set of hidden states. Each hidden can generate observable symbols given a specific emission function. Transition arcs allow the transition between two hidden states. There are thus two types of probability parameters: the state transition probabilities and observable symbol output probabilities. Given a finite sequence of hidden states, all the possible transition probabilities and symbol output probabilities can be multiplied at each transition to calculate the overall likelihood of all the output symbols produced in the transition path up to that point. Summing all such transition paths, one can then compute the likelihood that the sequence was generated by the HMM.

Figure 1 shows a graphical representation of a 3-state HMM, where a set of hidden states *S* has transition probabilities matrix *A* defined as a set of a_{ij}'s. The model is said to respect the first order Markov assumption if the transition from the current state to the next state only depends on the current location. Formally, a first order Markov model satisfies: $a_{ij} = P(q_t = S_j \mid q_{t-1} = S_i) = P(q_t \mid q_{1:t-1})$. Each hidden state has a specific symbol emission probability density function over the set of observable states *O*, as denoted by $B=P(o|S_i)$ for all i. Finally, completing the full definition of the HMM are the initial states probabilities π. A full HMM model is thus defined as $\lambda = (Q, A, O, B, \pi)$.

Three main computational issues need to be addressed with HMMs. The first issue is *state estimation*, answering the equation: $P(q_t = S_i | O_1 O_2 \dots O_t)$, which describes the most likely hidden state at time t is given by a sequence of observables from time 1 to t. The state estimation problem is solved with the forward/backward dynamic programming algorithm. The second issue consists of determining the most probable path of hidden state given a sequence of observable: this is solved by the Viterbi algorithm. Finally the last problem is the learning of the model, which is, given a sequence of observable, what is the maximum likelihood HMM that could produce this string? The learning problem is commonly solved by using an Expectation-Maximization (EM) method called the Baum-Welch algorithm (mathematical details of the different algorithms can be found in [17]).

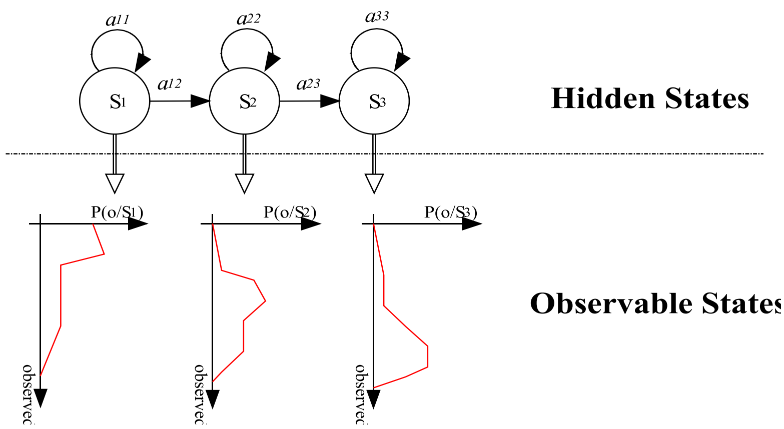

Fig. 1. A three-state Hidden Markov Model.

HMMs have been extremely successful in fields such as speech recognition [18], financial data [19], signal processing [20], and generic temporal data clustering [21]. There are, however, issues with HMMs in their nominal form. One problem is that HMMs rely on the assumption that the next state only depends on the current state (memory-less property) means that capturing higher level behaviors in the presence of noisy data can be difficult. Another issue inherent to the HMM structure lies in defining the basic structure of the model. Defining the model is a non-trivial issue, both in terms of number and meaning of states. While a simple model with too few states will not be able to model the complexity of the behavior properly, an excessive number of hidden states will results in an excess of parameters and possible over-fitting of the HMM, particularly if it is trained on scarce data set containing complex patterns. While there are no closed-loop solutions to determine the optimal number of states in a model, one possible solution is to use information theoretic metrics such as the Bayesian Information Criterion (BIC) or the Akaike Information Criterion (AIC), both of which measure $P(o | \lambda)$, the likelihood of the data given a model, while including a penalty for model with a large number of states [22].

In addition, while the parameter adjustments during the training of HMMs is a well researched problem, the states' emission functions and transition matrix of the trained model may or may not reflect the expected data distribution due to noise in the training data or inappropriate initial values for the model. HMM parameters, therefore, provide indirect semantic information about the training data [23] and the resulting model must be interpreted in order to highlight explanatory mechanisms.

Despite the issues highlighted above, we posit that the doubly stochastic nature of the HMMs is well suited to modeling human behavior: the cognitive states of an operator are not directly visible, but can be inferred from his or her actions. This is similar to the HMM notion of hidden states that must be inferred from observable symbols. Thus, we liken the hidden states of the HMM to the unobservable cognitive states of an operator that must be inferred from directly observable behaviors such as user interface interactions. Based on similar premises, some researchers have used HMMs to model human attention allocation based on eye-tracking data [24]. In contrast, we use HMMs in order to detect patterns in user behavior based on user interface events as the observable symbols. The hidden states then can be seen as higher cognitive states or intent that gave rise to the observable actions. For example, a UV controller selecting a UV would be an observable symbol, whereas his intention to replan the UV would be a cognitive, hidden state.

4 Experimental Data and Methodology

4.1 Experimental Data

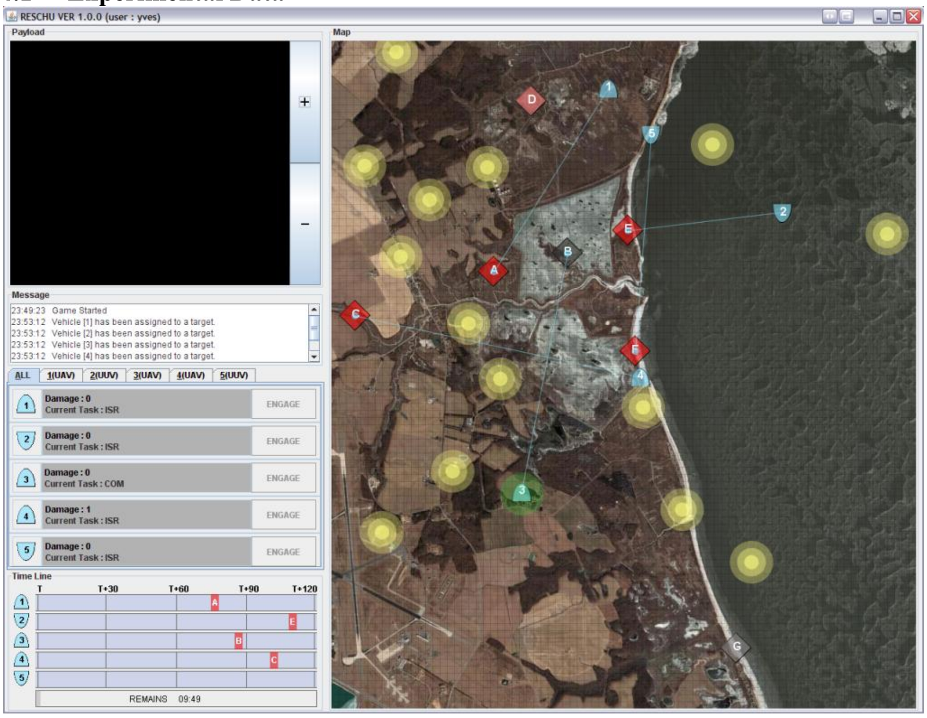

Fig. 2. The RESCHU interface

The data used in this work was obtained from the experiment described in [25]. While the goal of the original experiment was to validate a discrete event simulation of an operator controlling multiple heterogeneous unmanned vehicles, the recorded user interface interactions represents a rich corpus of supervisory control behavior. In the experiment, the Research Environment for Supervisory Control of Heterogeneous Unmanned Vehicles (RESCHU) Simulator was used to allow single human operators to control a team of UVs composed of unmanned air and underwater vehicles (UAVs and UUVs) (Figure 2).

In RESCHU, the UVs perform surveillance tasks with the ultimate goal of locating specific objects of interest in urban coastal and inland settings. UAVs can be of two types: one that provides high level sensor coverage (High Altitude Long Endurance or HALEs, akin to a Global Hawk UAVs), while the other provides more low-level target surveillance and video gathering (Medium Altitude Long Endurance or MALEs, similar to a Predator UAVs). In contrast, UUVs are all of the same type. Thus, the single operator controls a heterogeneous team of UVs which may consist of up to three different types of platforms, each with different characteristics.

Both UUVs and MALEs allow the user to perform a visual target acquisition. The visual task consists of looking for a particular item in an image by panning and zooming the camera view. The HALEs process emergent non-identified targets. The operator does not do a visual identification with HALEs. Once a HALE identifies a target, it is visually processed by either MALEs or UUVs. Once a UV has processed a target, an automated planner chooses at random the next target assignment, thus creating possibly sub-optimal target assignments that the subject can correct. Furthermore, threat areas appear dynamically on the map, and entering such an area could damage the UV, so the subject can optimize the path of the UVs by assigning a different goal to a UV or by adding waypoints to a UV path in order to avoid threat areas.

Participants maximized their score by 1) avoiding threat areas that dynamically changed, 2) completing as many of the visual tasks correctly, 3) taking advantage of re-planning when possible to minimize vehicle travel times between targets, and 4) ensuring a vehicle was always assigned to a target whenever possible.

Training was done through an interactive tutorial and an open-ended practice session. After participants felt comfortable with the task and the interface, they started the ten minute experimental session. After completing the experiment, the participants could see their score which corresponded to the total number of targets correctly identified. All data were recorded to an online database. The logged data of interest for this project consisted of user interactions with the interface in the manner of clicks, such as UV selections on the map or on the left sidebar, waypoint operations (add, move, delete), goal changes and visual tasks start and end.

4.2 Hidden Markov Models of Human Behavior

The creation of Hidden Markov Models of user behavior requires training the model on behavioral data. The raw behavioral data, which consists of the logged user interface events, cannot be used directly by the HMM learning algorithms, and must be pre-processed. Figure 3 shows the information flow of this process, which consists of a grammatical and a statistical phase.

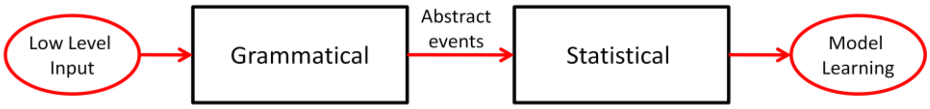

Fig. 3. A combined grammatical and statistical approach is used to infer future behavior from a stream of current behavior.

First, in the grammatical phase, low level input is translated into abstract events by the use of a syntactic parser. The grammatical rules used are shown in Table 1. User interactions are first separated based on the type of UV being controlled. Then, the interactions with each of the UV types are separated into different modes: selection either on the sidebar or on the map, waypoint manipulation (addition, deletion and modification), goal changes, and finally, the visual task engagement. Each entry of the table is then denoted by an integer index. The sequences of low-level user interactions are thus translated into a sequence of integers suitable for the statistical phase.

Table 1. The RESCHU grammar.

All					
Underwater UV					
Aerial UV					
High Altitude UV					
UV Type / **Mode**	**Select Sidebar**	**Select Map**	**Waypoint**	**Goal**	**Visual Task/Engage**

The second step of the process is the statistical learning phase during which the HMM is exposed to the training sequences of discrete data. More specifically, the training consists of maximizing the likelihood that the sequence observed data was generated by the HMM. This is done by adjusting the parameters of the HMM, such as the state transition matrices and the emission probability functions. The Baum-Welch expectation maximization algorithm can be used for this task. Because we do not know *a priori* what the expected patterns, i.e. the hidden states, should be, we do not have access to sequences of labeled observable data. Therefore, we use the EM algorithm in an unsupervised mode which will use Bayesian inferences in order to automatically infer the optimal parameters of the model. The EM algorithm, however, essentially performs a gradient search in parameters space to optimize the likelihood of the data given a specific model. It is, therefore, very sensitive to initial values. Furthermore, the structure, the number of hidden states, of the model must be provided for the algorithm to perform. A solution to both problems is to train multiple models of different length and with varying initial parameters.

5 Results

5.1 Selecting a Model Size

The experimental data consists of 25 test subjects, whose data sets were separated between 20 training sequences and 5 test sequences. The model training and testing was replicated with a different set of 5 test sequences to alleviate possible test set biases. The results reported here correspond to the averages of both training sessions. The test sets were used to validate the model and verify that the HMM does not overfit the relatively small training data corpus. Each model was trained until the probability of the training data converged.

After training multiple models of different sizes and with different starting points, we realized that, unless the initial probability density functions of the model were set to be equiprobable, most models of the same length tended to converge towards the same values regardless of the initial conditions. We then explored different possible model sizes and validated their performance on a testing set.

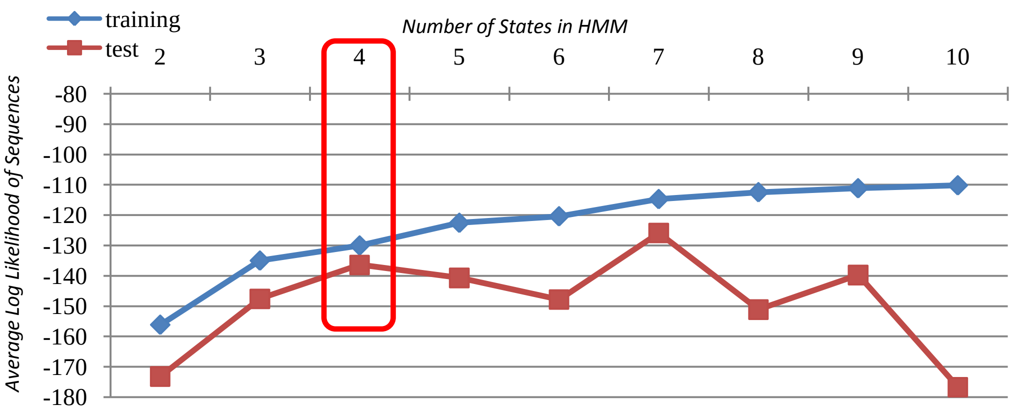

Fig. 4. Average log likelihood of the training and test sequences

Figure 4 shows the average log likelihood of the training and the test sequences. Note that the log likelihood scale is used in order to avoid very small numerical probabilistic values due to the length of the sequences. Figure 4 shows that the log likelihood of the training data increases as the number of states increases, which is consistent with the notion that the higher number of states allow for a closer fit of the training data. However, the log likelihood of the testing set starts to decline after 4 states and after 7 drops quickly. We therefore conclude that a 4 state model is a good fit for the observed data (for both the testing and training set).

5.2 HMM of an Operator Controlling Multiple Heterogeneous UVs

The parameters of the HMM allow for the interpretation of the structure of the underlying process. Figure 5 shows the transitions between the 4 hidden states of the

HMM obtained in the process described above. The label of each state is extracted from the state's probability density function over the space of observables. State 0, in this case, corresponds to behaviors consistent with planning and engaging UUVs and UAVs (including waypoint manipulation, goal selections and visual tasks), while State 3 corresponds to a monitoring behavior for UUVs and UAVs (selecting UVs on the map or on the sidebar). State 1 and 2 correspond essentially to the same states but for the HALEs. This symmetry is interesting for two reasons. The first is that it reveals the two main modes of behavior of the human operator during this task: monitoring and planning/engaging. This result is consistent with the human supervisory control literature [26]. Looking at the transition probability between the planning/engaging states and the monitoring states, we see that there is a clear tendency to successively alternate between the two behaviors (about 70% of the time, on average). The second conclusion we can draw from this symmetry is that the HMM was able to group UVs that possess the same functional characteristics: UUVs and MALEs were grouped by the algorithm together whereas the interactions with HALEs were seen as a distinct behavior. This result is explained by the fact that HALEs interactions do not involve a visual task, whereas both UUVs and MALEs do.

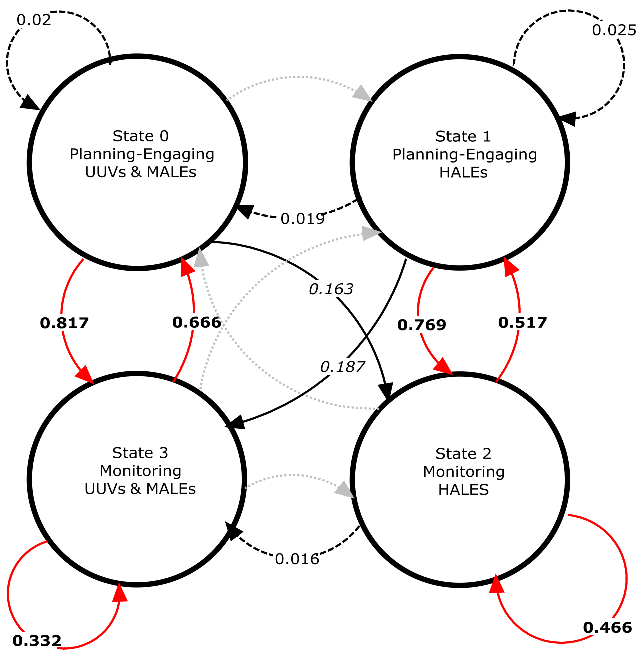

Fig. 5. The 4-state model of an operator of multiple heterogeneous UV.

Examining the transition between the two functionally different types of UVs (that is UUVs and MALEs versus the HALEs), we notice that the interaction switches between different types of UV happen only after a planning/engaging phase (about

16% of the time). This means that 1) the people tend to keep their interactions focused on functionally identical groups of UVs, and 2) that a common behavior would be to engage a type of UV and then switch to monitoring other functionally similar UVs. Emergent events, however, sometime force the monitoring the alternative type of UV.

The structure of the HMM is inherently generative in the sense that, given a specific state, one can generate the most likely set of state transitions along with the associated most likely sequence of observables. This means that we could predict what the next type of behaviors of the controller will be. The notion of "type of behavior" is important because the predictions are made at the level of the hidden states, which then drive the emission of observables. In practical terms, this means that the HMM would not be able to replace human controllers, but would be able to monitor their behaviors and detect deviations from the expected norm. These deviations, in turn, could be of significance as they could indicate that the operators are exhibiting abnormal behaviors which could eventually lead to human error. The HMMs, then, could be used as a valuable monitoring tool for team supervisors.

6 Conclusion

This paper showed a method to generate Hidden Markov Models of human operators controlling multiple heterogeneous unmanned vehicles. The process consists of a grammatical phase needed to translate user interface events into data suitable for the HMM formalism. Then, training the model involves finding out the optimal model structure. In this work, we found that a 4-state HMM provided a good fit for the behavior of a human supervisory controller of heterogeneous unmanned systems. We showed that the obtained HMM was able to recognize behavioral patterns consistent with the observed behaviors. Furthermore, unexpected patterns such as the rare switches between UV types were also highlighted. Finally, we also discussed how the transition matrix of the HMM model could be used to predict the future behaviors of an operator.

Acknowledgments. This research was supported by Boeing Phantom Works and the Office of Naval Research. The authors also would like to thank the members of the Humans and Automation for their contribution to this work.

References

[1] C. Nehme, R. M. Kilgore, and M. L. Cummings, "Predicting the Impact of Heterogeneity on Unmanned-Vehicle Team Performance," presented at 52nd Annual Meeting of the Human Factors and Ergonomic Society, New York, NY, USA, 2008.

[2] DOD, "Unmanned systems roadmap," Office of the Secretary of Defense, Washington, D.C. 2007.

[3] P. J. Mitchell and M. L. Cummings, "Management of Multiple Dynamic Human Supervisory Control Tasks," presented at The 10th International Command and Control Research and Technology Symposium (ICCRTS), McLean, VA, 2005.

[4] N. G. Leveson, "A new accident model for engineering safer systems," Engineering Systems Division, Massachusetts Institute of Technology, Cambridge, MA 2003.

[5] H. Simon, "A Behavioral Model of Rational Choice," in *Models of Man*, 1957.

[6] J. P. Barthelemy and E. Mullet, "Choice basis: A model for multi-attribute preference," *British Journal of Mathematical and Statistical Psychology*, vol. 39, pp. 106-124, 1986.

[7] D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, vol. 47, pp. 263-292, 1979.

[8] M. Endsley, "The application of human factors to the development of expert systems for advanced cockpits," presented at Human Factors Society 31st Annual Meeting, Santa Monica, CA, 1987.

[9] B. G. Silverman, M. Johns, K. O'Brien, R. Weaver, and J. Cornwell, "Constructing Virtual Asymmetric Opponents from Data and Models in the Literature: Case of Crowd Rioting," presented at 11th Conference on Computer Generated Forces and Behavioral Representation, 2002.

[10] J. McCarthy, "Programs with Common Sense," presented at Teddington Conference on the Mechanization of Thought Processes, 1958.

[11] C. A. Miller, "Modeling human workload limitations on multiple UAV control," presented at 48th Annual Meeting of the Human Factors and Ergonomics Society, New Orleans, 2004.

[12] J. A. Maule and A. C. Edland, "The Effect of Time Pressure on Human Judgement and Decision Making," in *Cognitive Models and Explanations*, W. R. Crozier, R. Ranyard, and O. Svenson, Eds., 1997.

[13] D. G. Wayne, E. J. Bonnie, and E. A. Michael, "The precis of Project Ernestine or an overview of a validation of GOMS," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. Monterey, California, United States: ACM, 1992.

[14] J. Anderson, *Rules of the Mind*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1993.

[15] G. Klein, *Sources of Power: How People Make Decisions*. Cambridge, MA: The MIT Press, 1999.

[16] D. Acuna and V. Parada, "Modeling expert players' behavior through data mining," presented at XXV International Conference on The Chilean Computer Science Society, 2005.

[17] L. Rabiner and B. Juang, "An introduction to hidden Markov models," *ASSP Magazine, IEEE [see also IEEE Signal Processing Magazine]*, vol. 3, pp. 4-16, 1986.

[18] J.-T. Chien and S. Furui, "Predictive Hidden Markov Model Selection for Decision Tree State Tying " presented at Eurospeech 2003, Geneva, 2003.

[19] Y. Zhang, "Prediction of Financial Time Series with Hidden Markov Models," in *School of Computer Science*, vol. Master of Applied Science: Simon Fraser University, 2004.

[20] D. H. Kil and F. B. Shin, *Pattern recognition and prediction with applications to signal characterization*. Woodbury, N.Y.: AIP Press, 1996.

[21] C. Li and G. Biswas, "Finding Behavior Patterns from Temporal Data Using Hidden Markov Model based on Unsupervised Classification," presented at International ICSC Symposium on Advances in Intelligent Data Analysis (AIDA'99), Rochester, N.Y., USA, 1999.

[22] R. Katz, "On Some Criteria for Estimating the Order of a Markov Chain," *Technometrics*, vol. 23, pp. 243-249, 1981.

[23] A. H. Kam, T. K. Ann, E. H. Lung, Y. W. Yun, and W. Junxian, " Automated recognition of highly complex human behavior," *International Conference on Pattern Recognition ICPR 2004*, vol. 4, pp. 327-330, 2004.

[24] M. Hayashi, "Hidden Markov Models to Identify Pilot Instrument Scanning and Attention Patterns," presented at IEEE International Conference on Man and Cybernetics, 2003.

[25] B. Donmez, C. Nehme, M. L. Cummings, and P. de Jong, "Modeling Situational Awareness in a Multiple Unmanned Vehicle Supervisory Control Discrete Event Simulation," *Journal of Cognitive Engineering and Decision Making, Special Issue on Computational Models of Macrocognition (in review)*, 2008.

[26] T. B. Sheridan, *Telerobotics, Automation and Human Supervisory Control*. Cambridge, MA: The MIT Press, 1992.