

Using Discrete-Event Simulation to Model Situational Awareness of Unmanned-Vehicle Operators

Carl Nehme, Jacob W. Crandall, and M. L. Cummings
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139

As the paradigm of operators supervising multiple unmanned vehicles becomes increasingly realizable, the impact on operator situational awareness of such a paradigm shift becomes very important. Quantifying the effects of alternate team configurations and system designs in terms of their impact on situational awareness is currently expensive, requiring time-consuming user studies. This paper presents an alternate method by which to study the impact on situational awareness in multi-UV control using a discrete event simulation model. A by-product of correctly quantifying operator situational awareness is the ability to use data to more accurately predict metrics such as mission performance and operator utilization. The paper also presents results from a user case study that was used to validate the effectiveness of using the discrete event simulation model to capture the effects on situational awareness as the size of the unmanned vehicle team being supervised is varied.

INTRODUCTION

IN order to achieve the military's future goal of one operator controlling multiple unmanned vehicles (UVs), the human operator's responsibility will shift from manually controlling vehicles to managing vehicles at the supervisory control level. Human supervisory control consists of higher-level tasking initiated by the human but delegated to the automation onboard the unmanned vehicles (Sheridan, 1992). Reduced workload afforded by supervisory control has several ramifications. An important and desired implication is the increase in operator idle time, which can be used to allow operator supervision of multiple vehicles simultaneously, hence inverting the current many-to-one ratio of operators to vehicles.

The possibility of a paradigm inversion has led researchers to examine the capacity of single operators to supervise multiple unmanned vehicles (Cummings et al., 2007; Olsen & Wood, 2003; Ruff et al., 2002). Earlier models of operator interaction with multiple vehicles allowed for the estimation of upper-bounds for the number of vehicles that can be supervised. However, such models assumed a serial arrival of requests for interaction from the vehicles as well as efficient operators that respond to requests for interaction instantaneously.

Since these assumptions are difficult to meet, the effects of having parallel arrival of requests for interaction as well as imperfect operators need to be accounted for. The first of these, parallel arrival of requests, has been addressed in general in scheduling theory and more specifically in work by Mau and Dolan (2006).

This paper utilizes a queuing model (Nehme et al., in submission) to show how the effects of imperfect operator situational awareness (SA) can be estimated using a simulation-based approach.

BACKGROUND

Although there have been several models introduced in the literature (Crandall et al., 2005; Olsen & Goodrich, 2003) for capturing the temporal aspects of the operators supervising multiple unmanned vehicles, these models have generally lacked the capacity to capture the effects of degraded SA. One drawback to this earlier work is the lack of accounting for human interaction delays and decision making inefficiencies.

The concept of Wait Times (WT) is an additional critical variable that is needed when modeling human control of multiple vehicles (Cummings & Mitchell, in press; Cummings et al., 2007). Although it is possible for human beings to multi-task, humans act as serial processors in that they can only solve a single complex task at a time (Broadbent, 1958; Welford, 1952). Even though it is possible for operators to rapidly switch between cognitive tasks, any sequence of tasks requiring complex cognition will form a queue and consequently, wait times will build (Cummings & Mitchell, in press; Cummings et al., 2007). Wait times can occur when 1) a vehicle is neglected while the operator is busy interacting with another vehicle, and/or 2) when an operator requires re-orientation time while switching between vehicles, and/or 3) when a vehicle is neglected due to lack of operator situation awareness.

To determine the effect of wait times on an operator’s ability to supervise multiple unmanned vehicles, Cummings et al. (2007) conducted an experiment with a UAV simulation test bed, holding constant the number of vehicles a person controlled. Cummings et al. (2007) showed that by using the wait times measured from the experiments, more conservative predictions for the number of vehicles that can be supervised resulted. Under both low and high workload conditions, the wait times due to the loss of situation awareness dominated overall wait times. The rest of this paper will introduce a discrete event simulation model that allows for the capturing of these wait times, and for the capturing of the effects such wait times have on mission performance and operator utilization.

SIMULATION MODEL

In this section, a discrete-event simulation (DES) model (Nehme et al., in submission), presented in Figure 1, is introduced that includes a) a queuing model of the human operator supervising multiple UVs, and b) a component for the ability to measure overall system performance.

Overview

The DES model was constructed under the assumption that the operator is acting in a supervisory control mode and that the different vehicles in the team are highly autonomous. By this assumption, this model is to be used to model operators whose role is that of a mission manager and whose task is to increase the performance of the unmanned vehicle mission. Also, the vehicles being supervised should, for the most part, have a predetermined plan functioning independently of the human, requiring necessary operator interaction only for tasks that require human judgment and reasoning. The operator can, however, attend to a vehicle even if the vehicle has not generated a task that requires human judgment and reasoning by initiating a re-plan that could potentially lead to improved performance. For example, in the case of an unmanned aerial vehicle that is assigned a reconnaissance task, the operator could re-plan the automation-generated vehicle path to better meet timing




Figure 1. DES Model for Heterogeneous UV System




Figure 2. Queuing Model

constraints.

The operator model is based on the single server queue architecture (Figure 2). An event in this queuing model represents tasks to be attended to by the operator. There are two main event types that are included in the model; a) events generated by the vehicles due to the arising of tasks that require human judgment and reasoning, and b) events that the operator induces by deciding to re-plan a vehicle’s existing plan.

The operator can attend to only one event at a time and this is captured by the single-server architecture; any events that arrive while the operator is busy will have to wait in the queue. This model allows for the capturing of operator interaction with the UVs through the servicing of those events that are arriving to the queuing system.

In this model, the arrival rate of vehicle-generated events is both a function of the rate at which events are generated by the vehicles as well as the timeliness by which they are noticed by the operator (which is a function of SA). The rest of this section will focus on this aspect of the operator model. For the definition of service times and for an explanation of the performance model, please refer to the complete model description presented in Nehme et al. (Nehme et al., in submission).

Arrival Rate of Vehicle-Generated Events

There is one event stream per vehicle. Considering just one of these streams, stream i , a vehicle-generated event arrives to the system if vehicle i generates a task that requires operator judgment/reasoning and the operator notices vehicle i ’s request. Before continuing its mission and generating other events, a vehicle that has already generated an event must first wait for operator attention; no new events corresponding to a vehicle can be generated while an event associated with that same vehicle already exists in the system. Thus, inter-arrival times for vehicle-generated events are the times between the completion of service for an event and the arrival of the next event. These inter-arrival times are described by a random variable (Λ_i) which has a particular probability distribution. The distribution of Λ_i is a function of two main components; a) the distribution of the random variable that describes the time following a service that it takes a vehicle to generate an event and, b) operator loss of SA (Equation 1).

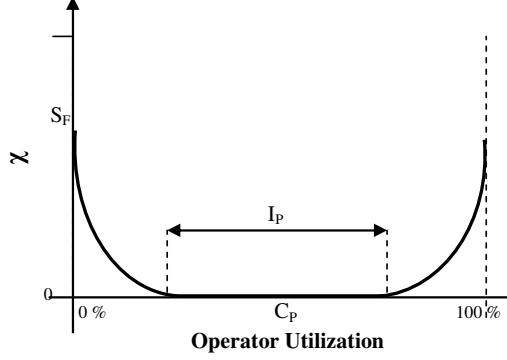


Figure 3. χ curve

$$\Lambda_i = X_i + \chi * X_i \quad (1)$$

The first term in Equation 1, X_i , is the random variable that describes the inter-arrival times between tasks that require operator judgment and reasoning. The probability distribution describing X_i is a function of the vehicle's autonomy, task complexity, and the environment of operation. Since the generation of a task by a vehicle does not necessarily imply that the operator notices the generated task, the first term excludes any effects due to loss of SA.

The second term in Equation 1, $\chi * X_i$, represents the penalty due to operator loss of SA, with χ taking a value of zero when the operator has complete SA and higher χ indicating degraded SA. The effect of low SA (i.e., high χ) is to create additional vehicle wait time (WTSA) which increases Λ , due to the operator taking longer to notice the vehicle (Cummings et al., 2007). The χ variable is multiplied by X to capture the dependence of the size of the penalties due to loss of SA on the rate at which the vehicle generates tasks that require operator intervention. Vehicles that produce tasks infrequently are serviced less often, and are therefore more likely to be overlooked than vehicles that are serviced more frequently.

Situational Awareness

SA is defined as the combination of perception of elements in the environment, the comprehension of their meaning, and the projection of their status in the future (Endsley, 1995).

To capture SA, this model builds on an assumption that SA is related to operator utilization (Endsley, 1993). When operators are under high levels of utilization, it is assumed that they are too busy to accumulate the information that is required to build SA. At the same time, when operators are under-utilized, it is presumed that due to a low level of arousal and complacency, they could overlook information from the environment, which would also lead to low SA.

In the DES, the χ variable in Equation 1 is related to operator utilization through a parabolic function that is concave upwards (Figure 3). This implies that at both high and low operator utilization, χ increases according to a quadratic

law and therefore increases Λ correspondingly. The parabolic relationship is inspired by the Yerkes Dodson Law (Yerkes & Dodson, 1908), which relates operator utilization to performance.

The size of the penalties at different utilization levels is dependent on the exact shape of the χ curve which is completely defined by providing a value for 3 different parameters. The first of these parameters is the point about which the curve is centered (C_p) which represents the utilization level where the SA penalty is zero. The second parameter, I_p , is the width of the interval around C_p which also has zero SA penalties (if I_p is zero, then the penalty is zero only at C_p utilization and the χ curve is a parabolic curve). Finally, the third parameter, S_f , is a scaling parameter that affects the magnitude of the penalties that can be incurred. In the case of a symmetric curve, χ takes on the value S_f at the 0% and 100% utilization levels.

Next, a user study was conducted and the results were compared to predictions from the DES model while varying the shape of the χ curve. The aim was to conduct a preliminary investigation into the usefulness of the SA model while varying its properties.

CASE STUDY

Predictions from the model were compared to results from an operator-in-the-loop experimental study simulating a search-and-rescue mission. The experimental study and model parameters are discussed below.

Software Test-bed

The human-UV team was assigned the task of removing as many objects as possible from the maze in an 8-minute time period. The objects were randomly spread throughout the maze in initially unknown locations. However, as each UV moved about the maze, more of the map was revealed to the participant. There were 22 possible objects to collect during a session.

Removing an object from the maze involved a three-step process. First, a UV moved to the location of the object in the maze (i.e., target designation, mission planning, path planning, and UV monitoring). Second, the UV collected the object (i.e., sensor analysis and scanning). In the real world, performing such an action might require the human operator to assist in identifying the object with video or laser data. To simulate this task, users were asked to identify a city on a map of the mainland United States using *Google Earth*-style software. Third, the UV carried the object out of the maze via one of two exits.

The human-UV interface was composed of a two-screen display (Figure 4). The maze was displayed on the left screen which contained the positions of the UVs and the currently visible objects in the maze. The right screen displayed the map that was used to locate the cities. When a user desired to

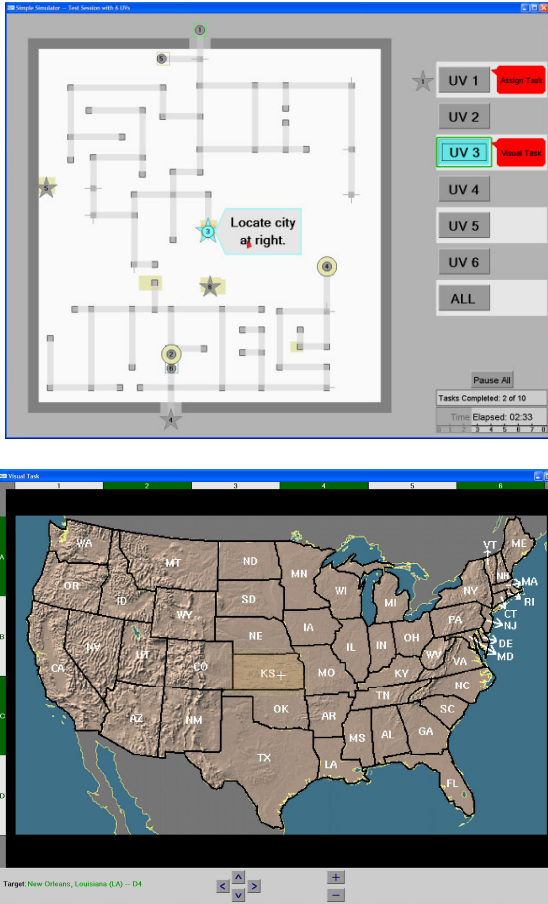


Figure 4. Two-screen interface by which an operator directed the UVs.

control a certain UV, s(he) clicked a button on the interface corresponding to that UV. The participant could then direct the UV by designating a goal location and/or modifying the UV's current trajectory.

Since the maze was not fully known, a UV had to choose between (a) moving along the shortest path of the known maze to its user-specified goal and (b) exploring the unknown portions of the maze in hopes of finding a shorter path.

Two different versions of UV autonomy were employed in the user study. In the first condition, called the *no-decision support* (NDS) condition, each UV's goal destination was determined completely by the human operator. Once the UV arrived at its user-defined goal destination, it did not move again until it received a new command from the user.

In the second condition, called the *full-decision support* (FDS) condition, each UV automatically selected a new goal when it was left idle (a management-by-exception level of automation was used). Additionally, if the user did not intervene, UVs automatically chose to exit the maze via the (estimated) nearest exit in the final 45 seconds of a session. The FDS condition also had one other additional decision

support tool to assist the user in locating cities on the map (to "pick up" objects). This decision support tool decreased the search time for a city on the map by about 5 seconds on average.

Experimental Procedure

This was a 2x4 mixed factorial study. The decision support conditions (NDS or FDS) was a between subjects factor. UV team size was a within subjects factor; each participant performed the search-and-rescue mission for team sizes of two, four, six, and eight UVs. The order in which the participants used each team size was randomized and counter-balanced throughout the study. After training on all aspects of the system, subjects then completed three comprehensive practice sessions. Following these practice sessions, each participant performed four test sessions (each with a different team size). Participants were paid \$10 per hour; the highest scorer also received a \$100 gift certificate. Thirty-two participants between the ages of 18 and 45 (mean 24.4 years old) participated in the study, 16 in each condition.

Model Parameters

To compare the DES model results to the human-on-the-loop experimental results, events in the case study corresponding to those described in the DES model were identified. Five data sets were measured from the experimental data: 1) arrival rate of vehicle-generated events after the UV was serviced for a vehicle-generated event, 2) arrival rate of vehicle-generated events after the UV was serviced for an operator-induced event, 3) service times of vehicle-generated events, 4) arrival rate of operator-induced events, and 5) service times of operator-induced events. The data sets collected were then used to generate random distributions that were used by the model.

The complete model of the human-UV team also requires a performance model. In the user study, the team scored points when an object was removed from the maze. In the NDS condition, this required two vehicle-generated events to occur (goal-assignment and locating a city). Thus, the DES Model awarded a point for the servicing of every two vehicle-generated events. In the NDS condition, only one vehicle-generated event (locating a city) needed to be performed. Thus, in this condition, the DES Model awarded a point for every serviced vehicle-generated event.

RESULTS

Using the random distributions generated from the service rates, re-planning rate and inter-arrival times, 10,000 trials were conducted using the DES model whose SA model was varied between four alternate configurations.

The first SA model was one where the penalty was zero under all levels of utilization ($I_p=100$) (Figure 5(a)). This is equivalent to ignoring SA in the DES model. The second

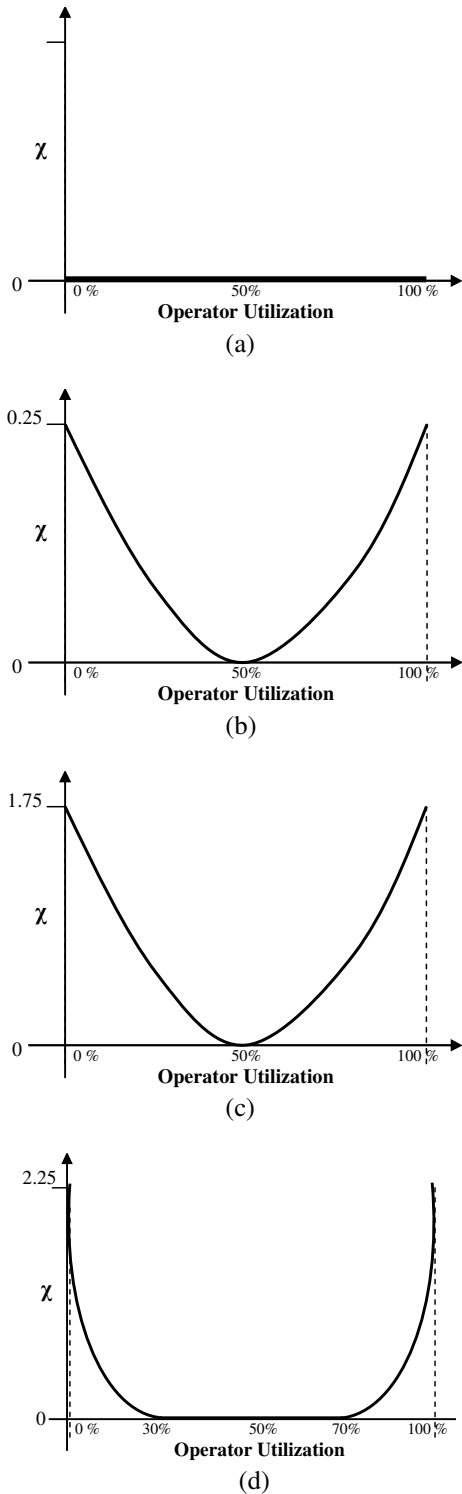


Figure 5. Four different situational awareness curve used to make predictions

model was a parabolic curve that was concave upwards where C_F was set to 50%, I_P set to 0, and S_F set to 0.25 (Figure 5(b)).

This is equivalent to using a simple parabolic curve with no magnitude scaling, as S_F is equal to 0.25 when scaling is ignored ($(1-0.5)^2$). The third model was also based on a parabolic curve where C_F was set to 50%, and I_P set to 0. However, in this case, the curve was scaled by setting S_F to 1.75 (Figure 5(c)). The magnitude scale factor for this parabolic curve was chosen by selecting the one that minimized the mean squared error for both performance score and operator utilization. Finally, the fourth model was one where C_F was set to 50%, I_P set to 40%, and S_F to 2.25 (Figure 5(d)). This is equivalent to using a curve where the penalty for loss of SA is zero when operator utilization is between 30% and 70% and with a magnitude scaling that minimizes the mean squared errors.

Effect of SA in Estimating System Performance

The actual scores in the user study for the full-decision support case are compared with results from the discrete event simulator for each of the four different SA curves in Figure 6 (similar profiles were found in the no-decision support case). The performance score mean squared error for each of the four SA curves was calculated (for both decision support cases). The SA model with no penalty performed the worst followed by the model with the simple parabolic curve. The two models with the scaling performed similarly with the model that had its interval of zero penalty modified performing slightly better.

It can therefore be concluded that the inclusion of the SA component in the model had an overall positive effect in terms of performance prediction accuracies, specifically when the model was scaled appropriately. When the effect of low SA (high or low utilizations) was completely ignored, the model tended to over-predict performance for the 4, 6 and 8 vehicle conditions and under predict for the 2 vehicle condition. This can be explained through the model overestimating/underestimating the number of tasks completed by the operator which led to higher/lower than expected performance scores. Including the SA effect, which was to increase the inter-arrival time between vehicle-generated events, led to decreased performance scores (improved predictions) except in the 2 vehicle condition where the model was already under predicting performance. Also, the improvement in the predictions when including the SA component depends on any scaling and shape modifications. Using a simple curve with no shape or magnitude modifications leads to a small improvement in the predictions while modifying the shape and magnitude parameters is required to make the penalties significant and the predictions more accurate.

Estimating Utilization

Average utilization was defined to be the ratio of the time the operator spends interacting with vehicles (servicing events), to total mission time. The inclusion of the SA component in the model also had an overall positive effect in terms of utilization predictions as can be seen in Figure 7.

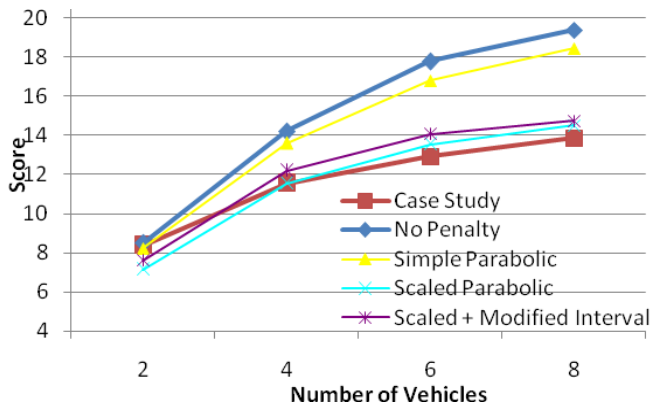


Figure 6. Effects of situational awareness model on mean scores for the full decision support case

When the effect of low SA (high or low utilizations) was completely ignored, the model tended to over-predict operator utilization for the 4, 6 and 8 vehicle conditions and under predict for the 2 vehicle condition. Just like in the case of the performance score, this can be explained through the model overestimating/underestimating the number of tasks processed by the operator which leads to higher/lower than expected utilization predictions. Including the SA effect led to an overall improvement in utilization predictions. This can also be explained by noting that the SA model tends to have a regulating effect on the number of tasks processed by the operator. When the operator is processing too many tasks, utilization is high which in turn leads to low SA (as computed by the model). The low utilization translates into increased penalties which in turn lead to less tasks being processed by the operator. When the SA model is excluded from the DES model, the lack of the regulation effect leads to overly large utilization estimates. As was the case for mean performance predictions, alternate SA curves lead to different results in terms of improved predictions.

CONCLUSIONS AND FUTURE WORK

The queuing model was powerful in modeling the human operator, particularly because it models both the impact of workload and situation awareness. While a critical aspect of human performance, situation awareness has been notoriously difficult to model. The queuing approach demonstrates one way its affects can be quantified. Moreover, by matching the SA sub-model to the results observed from experimentation with human subjects, it is possible to use the DES model to understand human behavior.

Since the DES model allows for alternative configurations in vehicle assignments, it can be used to understand the effect on SA as team composition is varied. Also, since the model

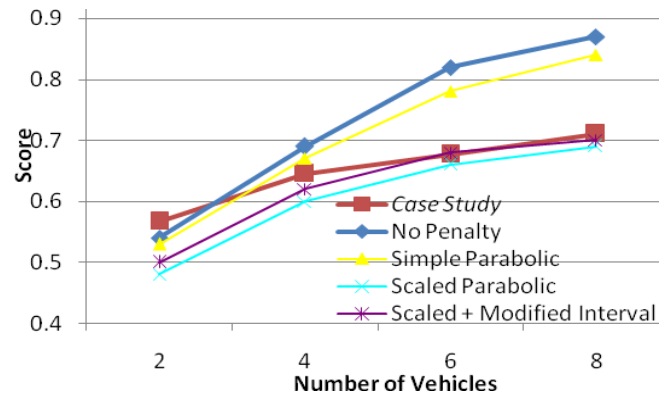


Figure 7. Effects of situational awareness model on mean utilizations for the full decision support case

allows for alternative operator strategies, the effect on SA can be studied while assuming different operator strategies. The operator strategies that have the least negative effects on SA can be designed for by modifying the decision support aides provided to the operator.

ACKNOWLEDGEMENTS

The research was supported by Charles River Analytics, the Office of Naval Research (ONR), and MIT Lincoln Laboratory.

REFERENCES

- Broadbent, D. E. (1958). *Perception and Communication*. Oxford: Pergamon.
- Crandall, J. W., Goodrich, M. A., Olsen, D. R., & Nielsen, C. W. (2005). Validating Human-Robot Interaction Schemes in Multitasking Environments. *Systems, Man, and Cybernetics*, 35(4), 438-449.
- Cummings, M. L., & Mitchell, P. J. (in press). Predicting Controller Capacity in Remote Supervision of Multiple Unmanned Vehicles. *accepted, IEEE Systems, Man, and Cybernetics, Part A Systems and Humans*.
- Cummings, M. L., Nehme, C. E., Crandall, J., & Mitchell, P. J. (2007). Predicting Operator Capacity for Supervisory Control of Multiple UAVs. In L. Jain (Ed.), *Innovations in Intelligent UAVs: Theory and Applications*.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32-64.
- Endsley, M. R. (1993). *Situation Awareness and Workload: Flip Sides of the Same Coin*. Paper presented at the

- 7th International Symposium on Aviation Psychology, Columbus, Ohio.
- Mau, S., & Dolan, J. M. (2006). *Scheduling to Minimize Downtime in Human-Multirobot Supervisory Control*. Paper presented at the International Workshop on Planning and Scheduling for Space.
- Nehme, C. E., Mekdeci, B., Crandall, J., & Cummings, M. L. (in submission). The Impact of Heterogeneity on Operator Performance in Futuristic Unmanned Vehicle Systems. *The International C2 Journal*, 1(3).
- Olsen, D. R., & Goodrich, M. A. (2003). *Metrics for Evaluating Human-Robot Interactions*. Paper presented at the Performance Metrics for Intelligent Systems.
- Olsen, D. R., & Wood, S. B. (2003). *Fan-out: Measuring Human Control of Multiple Robots*. Paper presented at the CHI2004.
- Ruff, H. A., Narayanan, S., & Draper, M. H. (2002). Human Interaction with Levels of Automation and Decision-Aid Fidelity in the Supervisory Control of Multiple Simulated Unmanned Air Vehicles. *Presence*, 11(4), 335-351.
- Sheridan, T. B. (1992). *Telerobotics, Automation, and Human Supervisory Control* Cambridge, MA: MIT Press.
- Welford, A. T. (1952). The psychological refractory period and the timing of high-speed performance - a review and a theory. *British Journal of Psychology*, 43, 2-19.
- Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology*, 18, 459-482.