AIAA <i>Infotech@Aerospace</i> 2007 Conference and Exhibit
7 - 10 May 2007, Rohnert Park, California

AIAA 2007-2704

Attention Allocation Efficiency in Human-UV Teams

Jacob W. Crandall* and M. L. Cummings†

Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge MA 02139

Human operators in unmanned vehicle systems must divide their attention between multiple tasks. This is particularly true of systems that contain multiple unmanned vehicles. In such systems, long-term research goals include determining (a) what dictates how humans allocate their attention, (b) where humans *should* focus their attention, (c) what can system designers do, via decision support systems, etc., to direct their attention to the "right" places, and (d) how do we measure the efficiency of human attention allocation. As a step toward finding answers to these questions, we conducted a user study in which a single human controlled multiple simulated unmanned vehicles. In this paper, we discuss some of the lessons we learned from this user study with respect to human attention allocation. These conclusions are made by using metrics derived from users' *selection strategies* and *switching times*.

I. Introduction

When a human interacts with multiple unmanned vehicles (UVs), (s)he must divide his/her attention between various tasks, including the various UVs in the team. In large part, the efficiency of human attention allocation determines the effectiveness of these systems. Thus, it becomes essential that we can answer questions such as:

- What dictates how humans allocate their attention among the various tasks to be performed?
- Where should human operators focus their attention at any given time?
- What can be done, in the form of decision support systems, system autonomy, etc., to direct operator attention to the "right" places?
- How should the efficiency of human attention allocation be measured?

To this end, we conducted a user study in which a single human controlled multiple simulated UVs performing independent tasks. From this study, we investigated three separate topics related to the previously mentioned questions. First, do users group tasks of like type, and then service them in succession? Second, how does a visually presented list of UVs affect human attention allocation? Third, with respect to attention allocation, what distinguishes successful operators from less successful operators? We used measures derived from switching times and attention allocation policies (or selection strategies) to answer these questions.

The remainder of the paper proceeds as follows. In Section II, we review related work. In Section III, we formalize the use of various terms and concepts related to attention allocation. In Section IV, we describe the user study, including the UV test-bed and experimental procedure. In Sections V, we explore the operator attention allocation characteristics related to the three topics discussed in the previous paragraph. In Section VI, we summarize the results.

II. Related Work

Due to the large body of research that has been conducted in the area of human attention allocation in human-machine systems, we review briefly only a small subset of this research in this paper. We refer the reader to Wickens and Hollands book[1] for a more extensive review.

*Postdoctoral Associate, Department of Aeronautics and Astronautics, Cambridge MA 02139. AIAA Member
†Assistant Professor, Department of Aeronautics and Astronautics, Cambridge MA 02139. AIAA Senior Member

Attention allocation efficiency (AAE) has been identified as one of three critical aspects of an unmanned vehicle system (UVS).[2] In general, AAE is the degree to which the human operator allocates her/his attention to the "right" places at the "right" time. In the context of an unmanned vehicles system (UVS), this is akin to allocating attention to the correct UV. Furthermore, human attention toward a UV can be allocated among different aspects of the UV's mission, including control, navigation, and mission management.[3]

Assigning task priorities is a critical part of any discussion of attention allocation. Task prioritization involves (a) determining how humans prioritize tasks,[4,5] (b) determining the best (or good enough) prioritizations, and (c) determining how to help operators focus their attention on the tasks with the highest priority. In the case of the latter point, specific efforts to focus human attention on the right tasks include developing reliable alarming systems[6,7] and changing the level of automation[8,9] to help operators retain adequate levels of situation awareness.[10] Additionally, human performance is known to be highest for moderate levels of workload.[11] Previous research has identified that, in many supervisory control contexts such as air traffic control, human utilization should not exceed 70%.[12,13]

The ability to measure AAE is a critical component of designing an efficient UVS. These metrics include the frequency of task switching,[14] costs of task switching,[15,16] and wait times caused be lack of situation awareness.[3] Whether these metrics are adequate measures of attention allocation efficiency or whether other generalizable metrics are needed remains an open research question.

III. Attention Allocation in Human-UV Systems

We now formalize the process of human attention allocation as we use it throughout this paper for a single operator, multi-UV UVS. For simplicity, we assume that attention allocation occurs at certain *choice points*. We consider that these choice points occur when the human operator ceases to interact with a particular UV and must determine which UV to service next.[a] We are interested in two aspects of this decision: (1) the UV that the operator chooses to service and (2) the amount of time it takes for the human to make this decision. The former aspect depends on the operator's *selection strategy* (SS). We refer to the latter aspect as switching time (ST).

Formally, let H_t be the history of the human operator's past interactions at time t. Let s_t^i be the state of UV i at time t, and let $S_t = (s_t^1, \cdots, s_t^N)$ be the systems joint state at time t, where N is the number of UVs in the team. Then, the human operator's selection strategy at time t, denoted $SS(t, H_t, S_0, \cdots, S_t)$, is a random variable specifying a probability distribution over UVs. In some cases, it is sufficient to make a Markov assumption in which the random variables is defined simply by the time t, the current joint state S_t, and the previous interaction (i.e., only the most recent component of H_t).

The amount of time it takes for the human to make a selection (or switching time) depends on two kinds of events. First, it depends on the amount of time it takes for the operator to orient to the status of the UVs in the team, determine which UV to service, and carry out the necessary steps to select that UV. Second, in many cases, the human may determine that no UV needs servicing for a period of time. Thus, the operator monitors the UVs until (s)he believes that one of the UVs should be serviced. We consider monitoring time to also be part of switching time in this paper. We note that the random variable describing switching time (ST) can be defined similarly to SS, except that the distribution is over times rather than UVs.

Ideally, estimates of SS and ST can be constructed by observing the UVS. These estimates can then be used to construct various metrics to tell the story of human attention allocation in the UVS. To demonstrate how these processes can be used, we present a user study in which a single user controls multiple simulated UVs. We describe this user study in the next section. We then use the resulting estimates of SS and ST to evaluate human attention allocation in this UVS in Section V.

IV. User Study

In this section, we describe the experimental UV test bed used including the mission to be performed, the user interface, and the behaviors of the UVs. We also describe the experimental procedure and designations of UV states used to construct models of SS and ST.

[a]This assumption means that we do not consider preemptive scheduling which occurs when a human terminates an interaction prematurely in order to service another task.

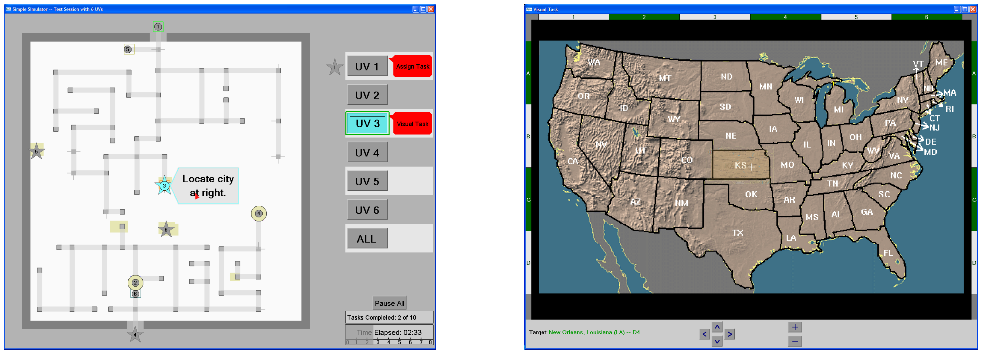

(a) (b)

Figure 1. The two displays used in the experiment. Each was displayed on a separate monitor.

IV.A. Mission

Participants used multiple simulated UVs to extract objects from a maze. The goal was (a) to remove as many objects from the maze as possible during an 8-minute session while (b) ensuring that all UVs were out of the maze when time expired. An object was removed from the maze using a three-step process. First, a UV moved to the location of the object in the maze. Second, the UV "picked up" the object. As this action would likely require the operator to perform a visual task, we simulated this task by asking the user to identify a city on a map of the United States using *Google*$^{\text{TM}}$*Earth*-style software; see Figure 1b. This locate-a-city task was a primary task and not a secondary task. Third, the UV carried the object out of the maze via one of two exits.

The objects were randomly spread throughout the maze. Initially, the positions of just six of the objects were shown. Every minute of the session, an additional two objects were displayed on the screen. Thus, there was a total of 22 possible objects for the UVS to collect.

IV.B. Interface

The human-UV interface used in the study was the two-screen display shown in Figure 1. On the left screen (see Figure 1a), the known portion of the maze was displayed along with the positions of the UVs and objects in the maze. The right screen (see Figure 1b) was used to locate the cities.

In this study, the user could control just one UV at a time. The user designated this UV by clicking a button on the interface corresponding to the desired UV (labeled UV1, UV2, etc.), as seen on the right side of Figure 1a. The user could then control the UV by specifying goal destinations and making path modifications. Goal designation was achieved by dragging the goal icon corresponding to the UV in question to the desired location. Once the UV received a goal command, it generated and displayed the path it intended to follow. The user was allowed to modify this path.

To "pick up" an object, the user was required to locate a specific city on a map of the mainland United States once the UV reached an object. This was done using the interface displayed in Figure 1b (the right screen). Once the user located the state where the city was located, the user could zoom into the map of this state using the mouse. The user then searched the map to locate the correct city.

IV.C. UV Behavior

The maze was not known at the start of the session by the user or the UVs, but the UVs created and shared a map of the maze as they moved about it. The incompleteness of the map meant that each UV needed to balance two objectives while moving in the maze: (1) taking the shortest path to its goal along the known maze and (2) exploring the unknown maze in search of a shorter path. As a heuristic, a UV explored unknown paths if it thought those paths could possibly be shorter than the shortest known path to its user-specified goal. Users often felt that the resulting UV behavior was undesirable. The user was allowed to modify the UV's path if desired.

IV.D. Experimental Procedure

After being trained on all aspects of the system and completing a comprehensive practice session, each user participated in six 8-minute sessions. UV teams with two, four, six, and eight UVs were tested. In each of the first four sessions, a different number of UVs were allocated to the team. In the last two sessions, the team sizes from the first two sessions were repeated. Thus, 18 samples were taken for each UV team size. The conditions of the study were randomized and counter-balanced. The participants were paid $10 per hour; the highest scorer also received a $100 gift certificate. Twelve people between the ages of 19 and 44 years old participated in the study. The mean age was 27.5 years old.

IV.E. UV States

We assume throughout this paper that the UVs were always in one of three states, which we call states O, V, and M. A UV was considered to be in state O if it was *outside* the maze. Likewise, a UV was in state V if it was on on object waiting for human assistance in performing the locate-a-city (or *v*isual) task. UVs were also considered to be in states O and V if they were within 5 seconds of exiting the maze or arriving at the location of their user-designated objects, respectively. Otherwise, a UV was considered to be in state M (i.e., moving about in the *maze*).

V. Results

We now discuss the attention allocation characteristics exhibited by the users in the study. We explore three different questions. First, did users service UVs that were in like states in succession? Second, how did users respond to a visually displayed list of UVs? Third, what distinguished more successful operators from less successful operators? We explore each of these questions separately.

V.A. Did users service UVs in like states in succession?

In complex situations such as control of multiple UVs, humans often form complexity mitigation strategies. One possible strategy is for the operator to perform tasks that are easiest or most important first. This would result in users performing tasks of like type in succession.

We first give some formal definitions. Let $Pr(X|X)$ denote the probability that the users chose to service a UV in state X, given that (a) there was at least one UV in each of the three states and (b) the previous interaction was with a UV that began in state X. Likewise, let $Pr(X)$ be the probability that the users chose to service a UV in state Y (where $Y \neq X$) given that there was a UV in each of the three states.

Thus, if there is a tendency for users to perform tasks of like type in succession, we would expect $P(X|X)$ to be higher than $P(X)$. Figure 2 shows these probabilities for 4-, 6-, and 8-UV teams.[b] All team sizes show the trend that, on average, a UV is more likely to be serviced if the user just completed interacting with a UV that was in the same state. The differences between the probabilities for V and M dependencies are both statistically significant (as given by Mann-Whitney U-tests; $p = 0.038$ and $p < 0.001$, respectively; $\alpha = 0.05$), but the difference between the probabilities for the O dependency is only marginally significant ($p = 0.038$).

V.B. How did users respond to a visually displayed list of UVs?

On the right side of the main interface screen (see Figure 1a), a static list of UVs in the team was displayed to the user, as shown in Figure 3a. It was from this list, or *visual stack*, that users selected the UV they desired to service (by clicking on that UV's labeled button). The selected UV was highlighted in powder blue. Status warnings (if any) associated with each UV were displayed to the right of the UV label. Star icons to the left of the UV label indicated that the UV was outside of the maze.

Our question is, how did this visual stack affect selection strategies? The average number of times UV 1 (the top UV) was serviced as opposed to UV N (the bottom UV) is shown in Figure 4a. For each team size, UV 1 was serviced more often than UV N. Except for the 4-UV case, these differences are all statistically significant; see Table 1. The figure also shows the trend that the discrepancy between the number of times the top and bottom UV were serviced increased as the number of UVs in the team increased.

[b]We omit the 2-UV case as we were not able to gather sufficient data for this condition.

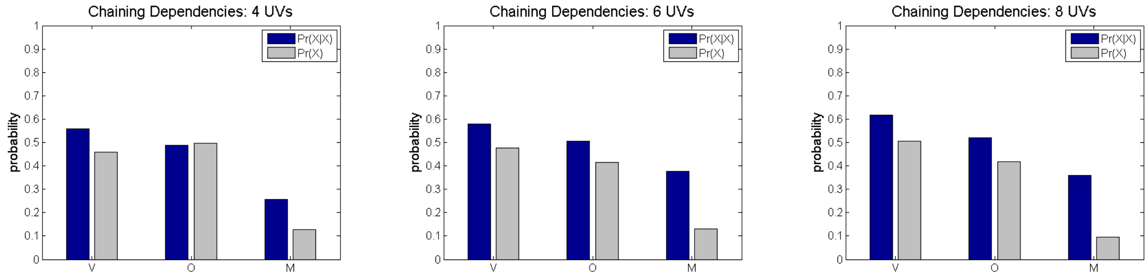

(a) (b) (c)

Figure 2. Influence of the previous service selection for 4-, 6-, and 8-UV teams.

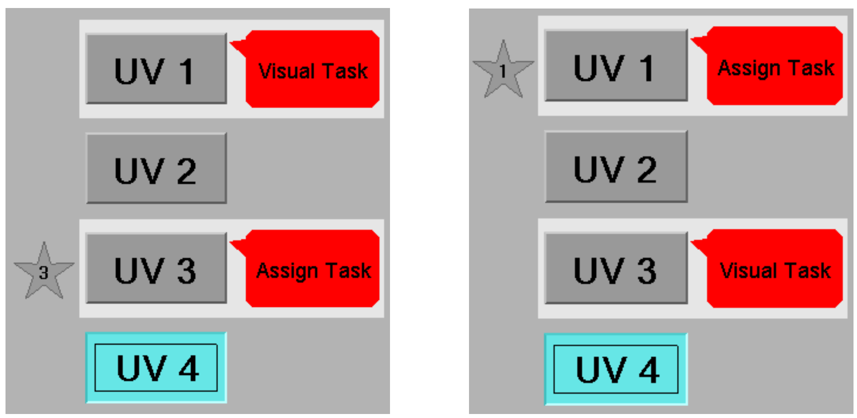

(a) (b)

Figure 3. A snapshot of the visual stack of UVs displayed on the main screen of the user interface for 2 different situations. (a) Situation in which a UV in state V is higher in the list than a UV in state O. (b) Situation in which a UV in state O is higher in the list than a UV in state V.

Likewise, when a user selected between multiple UVs in the same state, the UV in that state that was highest in the list was serviced more often than the others. For example, in the case multiple UVs in state V, the probability that the top UV in state V was serviced as opposed to random chance is shown in Figure 4b. For 4-UV teams and greater, the difference is statistically significant as determined by Mann-Whitney U-tests;[c] see Table 1. These results clearly show that users typically worked from the top of the stack, especially with larger team sizes.

While users typically worked from the top of the stack, we are still left with the question of whether the UVs states or the list order was more salient to the users. To answer this question, we consider situations like the ones depicted in Figure 3. In Figure 3a, UV 1 is in state V and UV 3 is in state O. Figure 3b shows the opposite situation in which UV 1 is in state O and UV 3 is in state V. If the UVs states were more salient for determining attention allocation than list order, then we would expect little difference between the probabilities that UV 1 would be selected in the cases depicted in Figure 3a and Figure 3b, respectively. However, we *would* expect a significant difference between these probabilities if list order was more salient than the UVs states.

To perform this test, we consider only the highest UV (in the list) in state V and the highest UV in state O given that (a) at least one UV is in state V and at least one UV is in state O and (b) one of the two UVs considered was selected. The average probabilities are shown in Figure 4c. In the figure, *V above O* indicates situations in which the UV in state V was higher in the list than the UV in state O, and *O above V*

[c]We used the Mann-Whitney U-test instead of a t-test when the data did not fit a normal distribution.

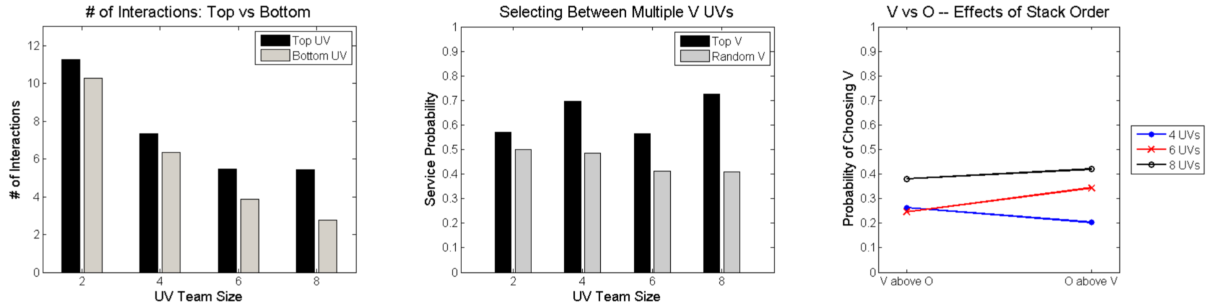

(a) (b) (c)

Figure 4. Changes in selection strategies based on order in the stack. (a) Contrasts the number of interactions with the top UV on the stack with the last UV. (b) Contrasts the probability that the top UV in the list in state V was serviced as opposed to random chance. (c) Shows the effects of list order on service probability for UVs in different states.

Team Size	T-test Top vs. Bottom UV	Mann-Whitney U-test Top V vs. Random
2	$t(17) = 2.43$, $p = 0.027$	$p = 0.710$
4	$t(17) = 1.52$, $p = 0.146$	$p = 0.022$
6	$t(16) = 3.27$, $p = 0.005$	$p = 0.003$
8	$t(17) = 4.09$, $p < 0.001$	$p < 0.001$

Table 1. Statistical reports for various tests referenced in the text.

indicates situations in which the UV in state O was higher in the list than the UV in state V. There are no statistical differences between the two probabilities of selecting the V state for any team size. Thus, the UVs states had a stronger effect on human attention allocation than did the order of the UVs in the list.

It is possible that the star icons to the left of the UV labels (indicating UVs in state O) had an impact on this result, meaning that visually salient star could have caused subjects to select the O state more often than the V states. However, we cannot determine this with our current data. Thus, we leave this analysis to future work.

Together, these results show that, while the order of the list impacted human attention allocation, the states of the UVs had a greater effect on human attention allocation than list order.

V.C. What distinguished more successful operators from less successful ones?

In our study, the success of the UVS was measured by the number of objects collected and the number of UVs lost (i.e., left in the maze when time expired). These variables varied widely from user to user. As different factors affect these two performance variables, we focus only on objects collected.

While the previous questions related only to selection strategies, this question relates to both selection strategies and switching times. We analyze both aspects in the form of two questions. First, did operators with a propensity to focus attention on UV progress and behaviors collect more or less objects than those that did not? Second, did operators with shorter switching times collect more objects than those with longer switching times? We address each question separately.

V.C.1. Should operators focus on UV behaviors?

As we mentioned previously, each UV used a simple heuristic to balance exploring the maze and moving toward its goal destination. The resulting behavior was sometimes opposite to users' intuitions and desires. This often led users to monitor and correct UV behaviors. Additionally, since the locations of more objects become known at random times throughout the mission, users often chose to reassign the UVs' goal destinations once they were already moving in the maze. We refer to users with tendencies to correct or reassign UV goals as *micromanagers*. In so doing, we do not mean to associate so-called micromanagement with a negative meaning, however, as some goal and path corrections were beneficial.

Team Size	Correlations Objects vs. Micromanagement
2 UVs	$r(16) = 0.34$, $p = 0.17$
4 UVs	$r(16) = -0.73$, $p < 0.001$
6 UVs	$r(15) = -0.19$, $p = 0.46$
8 UVs	$r(16) = -0.49$, $p = 0.038$

(a)

Team Size	Correlations Objects vs. Switching Times
2 UVs	$r(16) = -0.30$, $p = 0.230$
4 UVs	$r(16) = -0.85$, $p < 0.001$
6 UVs	$r(15) = -0.67$, $p = 0.003$
8 UVs	$r(16) = -0.69$, $p = 0.002$

(b)

Table 2. (a) Correlations between the number of objects collected and degree of micromanagement. (b) Correlations between the number of objects collected and average switching time.

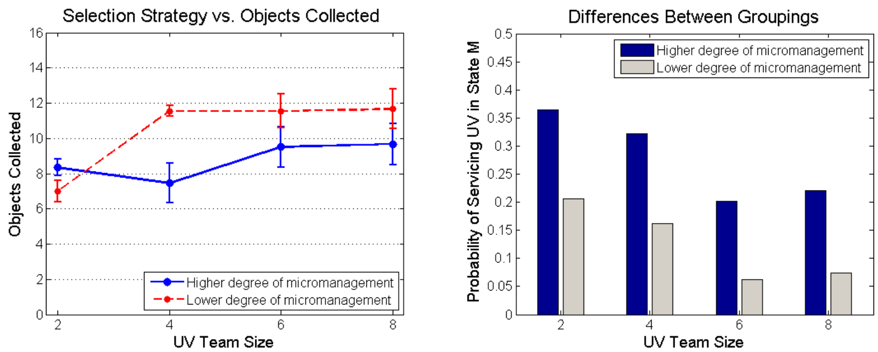

(a) (b)

Figure 5. Comparisons of UVS effectiveness based on high and low degrees of micromanagement. (a) A comparison of number of objects collected. (b) A comparison of the mean probability of servicing a UV in state M.

We defined a user's degree of micromanagement as the average probability that (s)he chose a UV in state M at a choice point. Correlations between degree of micromanagement and objects collected are given in Table 2a. The table shows a (significant) strong negative correlation between degree of micromanagement and objects collected for the 4-UV case, and a moderately negative (significant) correlation for the 8-UV case. However, the data show no real correlation in the 2-UV and 6-UV cases, although the sign change in the 2-UV case may indicate that some micromanagement in a low workload environment could be beneficial.

To better see the effects of micromanagement on UVS effectiveness, we divided users into two equally distributed groups based on their degree of micromanagement and plotted the mean performance of each group. The results are shown in Figure 5. The average number of objects collected by each group is shown in Figure 5a. For the 2-UV case, higher degrees of micromanagement produced higher scores. With larger team sizes, higher degrees of micromanagement produced lower scores.

A comparison of the mean degree of micromanagement of the two groups is shown in Figure 5b. This figure shows that the degree of micromanagement decreased on average (for both groups) as the size of the team increased. This trend is expected since the additional UVs represented increasing workload.

V.C.2. Do shorter switching times mean higher performance?

We now analyze how switching times correlate with the number of objects collected. Recall that switching times were the times between when an interaction with one UV ended and an interaction with another UV began. Thus, they included (a) the time spent reorienting to the status of the UVs in the team and selecting a UV, and (b) the time in which the user monitored the system (when (s)he believed that none of the UVs required servicing). To limit outside factors that could influence switching times, we consider only switching times in which at least one UV in the team was in state O or V.

Correlation results between the number of objects collected and length of switching times are given in Table 2b for all team sizes tested. The correlations are stronger than those for degrees of micromanagement; see Table 2a. The data show strong negative correlations between switching time and objects collected for the 4-, 6-, and 8-UV teams. There is no significant correlation for the 2-UV case.

The difference in objects collected between the group of users with lower switching times and the group

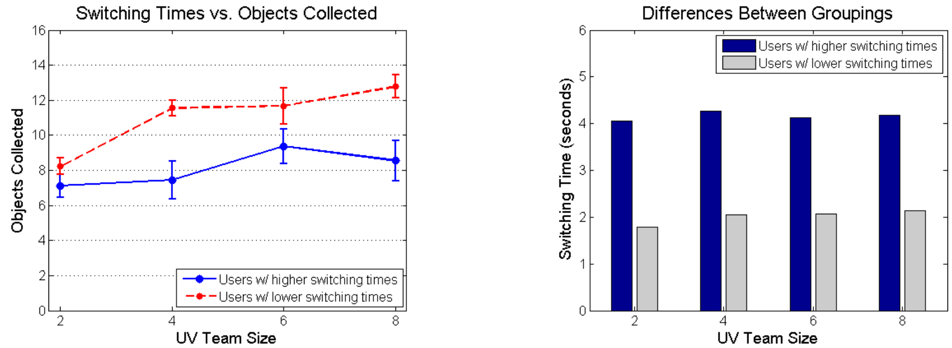

(a) (b)

Figure 6. Comparisons of UVS effectiveness based on high and low switching times. (a) A comparison of number of objects collected. (b) A comparison of mean switching times between the two groups.

of users with higher switching times (evenly divided) for the base system is shown in Figure 6a. The plot looks similar to the one obtained for groupings of micromanagers, though interesting differences can be seen in the 2- and 8-UV cases; see Figure 5a. The difference in mean switching times between the two groups is shown in Figure 6b.

It appears from Figure 5b that switching times are constant across UV team sizes. An ANOVA across UV team size verifies that this is, indeed, the case ($F(3, 67) = 0.01$, $p = 0.961$).

VI. Summary

In this paper, we discussed various aspects of human attention allocation in single operator, multi-unmanned vehicles systems. We presented a user study and used data from this study to answer questions about human attention allocation. Using metrics derived from operator selection strategies and switching times, we found that human operators had tendencies to attend to similar tasks in succession. We also found that attention allocation was influenced by the ordering of UVs in a visual list, but not to the extent that it interfered with attention allocation efficiency to a significant degree. Furthermore, we showed that high performance correlates with low switching times and, to a lesser extent, low degrees of micromanagement.

Acknowledgments

This research was funded by MIT Lincoln Laboratory.

References

[1]Wickens, C. and Hollands, J. G., *Engineering Psychology and Human Performance*, Prentice Hall, Upper Saddle River, NJ, 3rd ed., 2000.

[2]Crandall, J. W. and Cummings, M. L., "Developing Performance Metrics for the Supervisory Control of Multiple Robots," *Proceeding of the 2nd Annual Conference on Human-robot Interaction*, 2007, pp. 33–40.

[3]Cummings, M. L. and Mitchell, P. J., "Operator Scheduling Strategies in Supervisory Control of Multiple UAVs," *Aerospace Science and Technology*, 2007.

[4]Neth, H., Khemlani, S. S., Oppermann, B., and Gray, W. D., "Juggling Multiple Tasks: A rational Analysis of Multi-tasking in a Synthetic Task Environment," *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting*, San Fransisco, CA, 2006.

[5]Sheridan, T. B. and Tulga, M. K., "A Model for Dynamic Allocation of Human Attention Among Multiple Tasks," *14th Annual conference on Manual Control*, NASA Ames Research Center, 1978.

[6]Ho, A. W. L., Cummings, M. L., Wang, E., Tijerina, L., and Kochar, D. S., "Integrating Intelligent Driver Warning Systems: Effects of Multiple Alarams and Distraction on Driver Performance," *Transportation Research Board Annual Meeting*, 2006.

[7]Dixon, S. R., Wickens, C. D., and Chang, D., "Unmanned Aerial Vehicle Flight Control: False Alarms Versus Misses," *Proceedings of the Human Factors and Ergonomics Society 48th Annual Meeting*, 2004.

[8]Sheridan, T. B. and Verplank, W. L., "Human and computer control of undersea teleoperators," Tech. rep., Man-Machine Laboratory, Massachusetts Institute of Technology, Cambridge, MA, 1978.

[9]Parasuraman, R., Sheridan, T. B., and Wickens, C. D., "A Model of types and Levels of Human Interaction with Automation," *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, Vol. 30(3), 2000, pp. 286–297.

[10]Kaber, D. B. and Endsley, M. R., "The Effects of Level of Automation and Adaptive Automation on Human Performance, Situation Awareness and Workload in a Dynamic Control Task," *Theoretical Issues in Ergonomics Science*, Vol. 5(2), 2004, pp. 113–153.

[11]Yerkes, R. M. and Dodson, J. D., "The Relation of Strength of Stimulus to Rapidity of Habit-Formation," *Journal of Comparative Neurology and Psychology*, Vol. 18, 1908, pp. 459–482.

[12]Rouse, W. B., *Systems Engineering Models of Human-Machine Interaction*, New York: North Holland, 1983.

[13]Schmidt, D. K., "A Queuing Analysis of the Air Traffic Controller's Workload," *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 8(6), 1978, pp. 492–498.

[14]Wang, J. and Lewis, M., "Human control for cooperating robot teams," *Proceeding of the 2nd Annual Conference on Human-robot Interaction*, 2007, pp. 9–16.

[15]Squire, P., Trafton, G., and Parasuraman, R., "Human control of multiple unmanned vehicles: effects of interface type on execution and task switching times," *Proceeding of the 1st Annual Conference on Human-robot Interaction*, ACM Press, New York, NY, USA, 2006, pp. 26–32.

[16]Goodrich, M. A., Quigley, M., and Cosenzo, K., "Task Switching and Multi-Robot Teams," *Proceedings of the Third International Multi-Robot Systems Workshop*, 2005.