Manual control versus management-by-consent in managing multiple threats

Chris Baber[1] & Natan S. Morar[2]

School of Computer Science, University of Birmingham

ABSTRACT

As the use of Uninhabited Aerial Systems (UAS) increases, e.g., for commercial delivery, for surveillance or for hostile action, there are challenges of monitoring and appropriately responding to a crowded airspace. Providing automated support could reduce the challenges. However, such support might also have an impact on the strategy that a human operator deploys. In this paper we present a simulation of Air Defence (in the form of a single-player, interactive game) which is used to study human performance under three conditions. Provision of decision support, i.e., through management by consent, produced better performance, even though it provided a limited situation awareness. A hybrid display produces performance that is superior to the manual control condition and similar to the management by consent condition. We note that provision of the air picture alone resulted in a different form of suboptimal performance, in which sensitivity was significantly lower. In this respect, providing the decision support (in the form of the polygon display) helps to limit tendency for false alarms.

KEYWORDS

Air Defence, Multiple Threats, Management-by-Consent, Management-by-Exception, Manual Control

Introduction

This paper is motivated by the challenge of monitoring an air picture in which different types of Uninhabited Aerial Systems (UAS) could present different levels of threat. In this instance, we consider whether the provision of an automated decision aid improves an air defence operator's ability to respond to potentially threatening behaviour. Related to this is the question of whether it is necessary to include the air picture or whether operators can make decisions in response to the automated decision aid. One might expect the air picture (in the form of a graphical display of moving objects) to enhance Situation Awareness, but if this is cluttered then its benefit might be compromised. Alternatively, having a clear indication of a recommended action (from the decision support) could allow rapid and accurate response, but might provide limited Situation Awareness.

Previous work suggests that people can deal with up to 4 UASs with minimal automated support (Cummings et al., 2007; Liu et al., 2009; Baber et al., 2011). With automated support, people might be able to monitor up to 15 UASs (Porat et al., 2016), providing that the operator uses a composite display which provides clear information on status relating to specific categories of data, e.g., threat level or UAS status. A recommended approach to design of composite displays "…is to use a display with an emergent perceptual feature such as a polygon…" (Parasuraman et al., 2000, p.288).

The number of UAS interacts with the type of activity that the operator performs. Ruff et al. (2002) show that with manual control of individual units is challenging when the number of UAS is 4 or more. However, the number of UAS can be increased if there is a management-by-consent policy, i.e., automation suggests an action and the operator confirms (Ruff et al., 2002), They also found that management-by-exception, i.e., automation performs an action unless explicitly prevented, to be the least effective approach. Thus, in the present study, we present 15 UAS and response is made either in a manual mode, i.e., where the participant sees the UASs and needs to decide when to intervene, or management-by-consent, i.e., automation advises on when to intervene.

To summarize, we assume that:

- It is possible to automate air defence (particularly if there is a tightly specified region that corresponds to threat and tightly specified definition of threatening behaviour). In such circumstances, one could conceive of a fully automated system.
- However, one might be cautious about relying solely on a fully automated system because the specification of threat region or threatening behaviour might occasionally be sufficiently ambiguous to require redefinition (by a human operator). Further, the responsibility for response to a threat cannot be given to automation and so there is a role for 'human-over-the-loop' monitoring of such systems.
- In the absence of automation, human operators will rely on their Situation Awareness to determine appropriate responses to make to threats.
- However, as the number of drones increases, it might be difficult to fully maintain awareness and this could have an impact on performance (either resulting in the operator missing threats or in the operator being overly responsive, i.e., making more false alarms).

A Simulated Air Defence Task

We designed and built a simulated air defence task (the Drones Game, figure 1). The DroneGame is a simple, one-player game that emulates some aspects of Air Defence (Baber and Vance, 2020). A number of UASs (the game allows up to 20, but in this study we use 15) fly around an environment and the goal is to protect designated areas by disabling any UAS that represents a threat while allowing free passage to friendly UASs. As each UAS flies into range, a 'beacon' can be activated to reduce power in the UAS. Object classification is performed 'automatically', i.e., UAS are coloured red, yellow, blue or green to reflect threat level (i.e., hostile, unknown, neutral, friendly correspondingly). At set up, half of the UASs are classified as 'hostile or unknown' and half as 'neutral or friendly'. The groups are further divided to have similar numbers of each type (for odd integers, there will be one more hostile UAS in the first group and one more neutral in the second group). Of the 'hostile or unknown' group, half of the hostile UAS present as yellow and then change to red when they enter the range of the beacon.

Figure 1 shows the game screen, with UASs flying over houses in 'Los Angeles'. Houses that have a border around them are being protected from hostile UAS. In the instructions to participants, 'hostile intent' was defined in terms of airborne cameras being flown by paparazzi to take photographs of 'celebrities'. In order to disable a UAS, one selects a beacon (by clicking on the beacon or pressing the number key that corresponded to the beacon id) and then presses the space-bar to activate the beacon. Selecting the beacon causes the 'area of activation' (the circle around the beacon) to change colour (to purple) and decrease in size (as long as the beacon is activated). An activated beacon reduces power in any UAS within the area of activation. The aim is to ensure the hostile UAS have minimal power while allowing the friendly UAS to maintain maximal power.

If the UAS is disabled in the vicinity of the protected house, then this could cause damage and so the activation of a beacon should be performed when the UAS enters the beacon range (indicated by the green circle) and outside the area of risk (highlighted in orange). When the beacon is active, its range decreases (corresponding to use of energy). In order to recharge the beacon, it needs to be turned off (which requires the participant to select the beacon and use the space-bar to toggle it off).

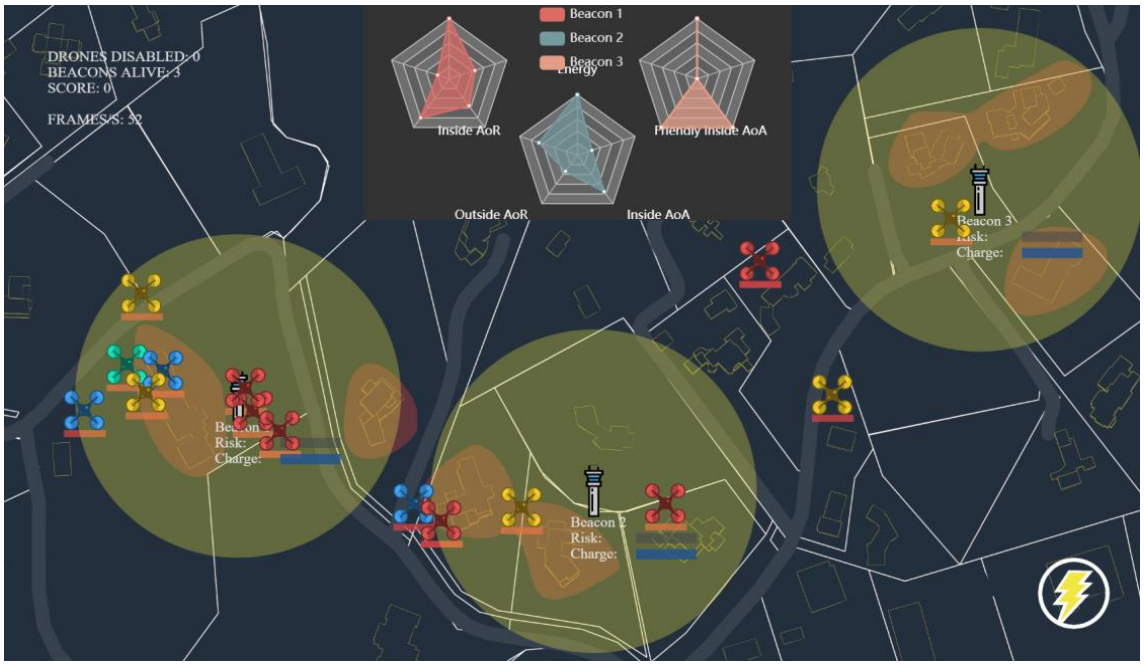

Figure 1: DroneGame screen

In the top, centre of the screen, there are three polygon displays (one for each beacon). These indicate the threat level and indicate when a response is required (figure 2).

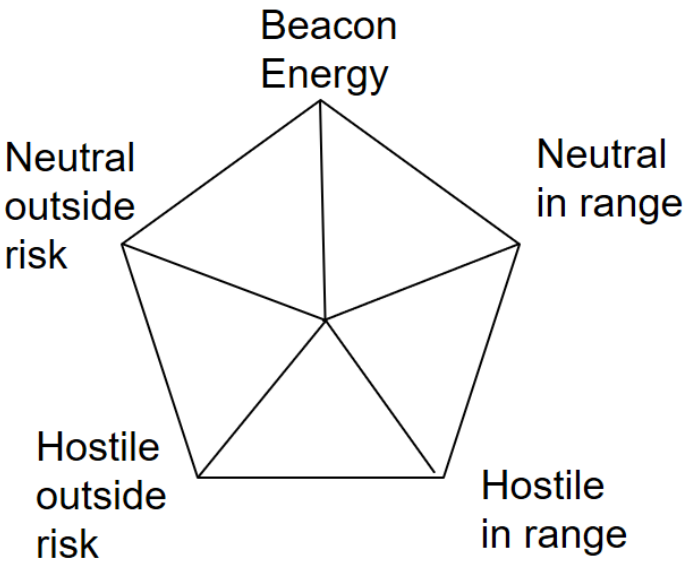

Figure 2: Polygon display

Experiment

The task is to monitor a dynamic air picture and disable hostile UAS. In the manual condition, the movement of the UAS is displayed in an air picture and the participant decides when to intervene. In management-by-consent conditions, automation indicates, using a polygon display, when a threat could be disabled. In one version of management-by-consent, participants only see the polygon display and in the other, they see both polygon display and air picture. We are interested in how these different conditions affect performance, in terms of response time and in terms of signal detection.

Participant Recruitment

A link to an online version of the Drones Games was circulated to 143 students on MSc Computer Science programmes at University of Birmingham. Ethical approval for the experiment was given by MODREC (914/MODREC/19) and University of Birmingham (ERN_18-1988). Participants had agreed to receive the link and were told that they could choose whether or not to play the game. They were required to install the game on their own computers and configure it according to instructions provided. 92 participants completed consent forms and attempted the game. The 'config' logs allowed the experimenters to check that configurations were correct. If a participant had made an error on any of the three configurations, or they had not logged data for any of the three conditions or had not activated the beacon in one or more of the conditions, their data were excluded. This resulted in a sample of 66 usable sets of data.

Task

All participants were presented with 15 drones and were asked to respond to both red and yellow drones (this was shown to be a more challenging task in previous experiments, i.e., Baber and Vance, 2020). Participants completed the game in all three conditions, i.e., manual, management-by-consent with computer advice only, management-by-consent with computer advice and air picture.

Data Preparation and Analysis

We report three measures of performance. Targeting performance is defined in terms of *Signal Detection d*' [calculated as, d' = ln(((H*(1-FA))/((1-H)*FA)))] in terms of UAS having their power reduced, where hits (H) are hostile UAS, and false alarms (FA) are either non-hostile targets (FA_target) *or* a drone (or either type) over an protected region. For initial calculations we combined both FA measures but report separate analysis of these in terms of strategy. *Time per Target* is the number of hits divided by the total time over which the game was played (typically around 10 minutes per condition). *Beacon Activation* is the count of the beacon being switched on or off. For each measure, the Data were tested for normality (using a Shapiro-Wilk test). If at least one condition failed this test, analyses were conducted using non-parametric tests. In all cases, Friedman tests were applied for Analysis of Variance, and Wilcoxon signed-rank tests applied for post-hoc pairwise comparison.

Results

A False Alarm (target) meant that the participant had hit a neutral or friendly UAS, and a False Alarm (region) meant hitting a UAS adjacent to protected house. Table 1 summarizes these results.

Table 1: Summary of Targeting performance [Mean (sd)] *p<0.05, **p<0.001

Condition	Hit	FA_target	FA_region
Management-by-Consent1 (Polygon display of computer advice only)	38(±47)	22(±45)	8 (±14)
Management-by-Consent2 (Polygon display plus air picture)	56(±54)	29(±47)	13 (±17)
Manual Control (air picture only)	68(±69)	38(±57)	17 (±20)
Friedman	X^2 =10.2**	X^2 =7.9*	X^2 =22.8**

In table 1, performance in the Manual condition has the highest rate of Hits. However, it also has higher rates for both FA measures. Sensitivity, d' (in terms of the relationship between Hits and False Alarm) shows that Manual Control had lower d' than the other conditions (figure 3). Average d' scores for Manual Control were significantly lower than Management-by-consent 1 ($z = 2.3$, $p<0.05$, power = 0.28) and Management-by-consent 2 ($z=2.1$, $p<0.04$, power = 0.26). There was no difference between the two management-by-consent conditions.

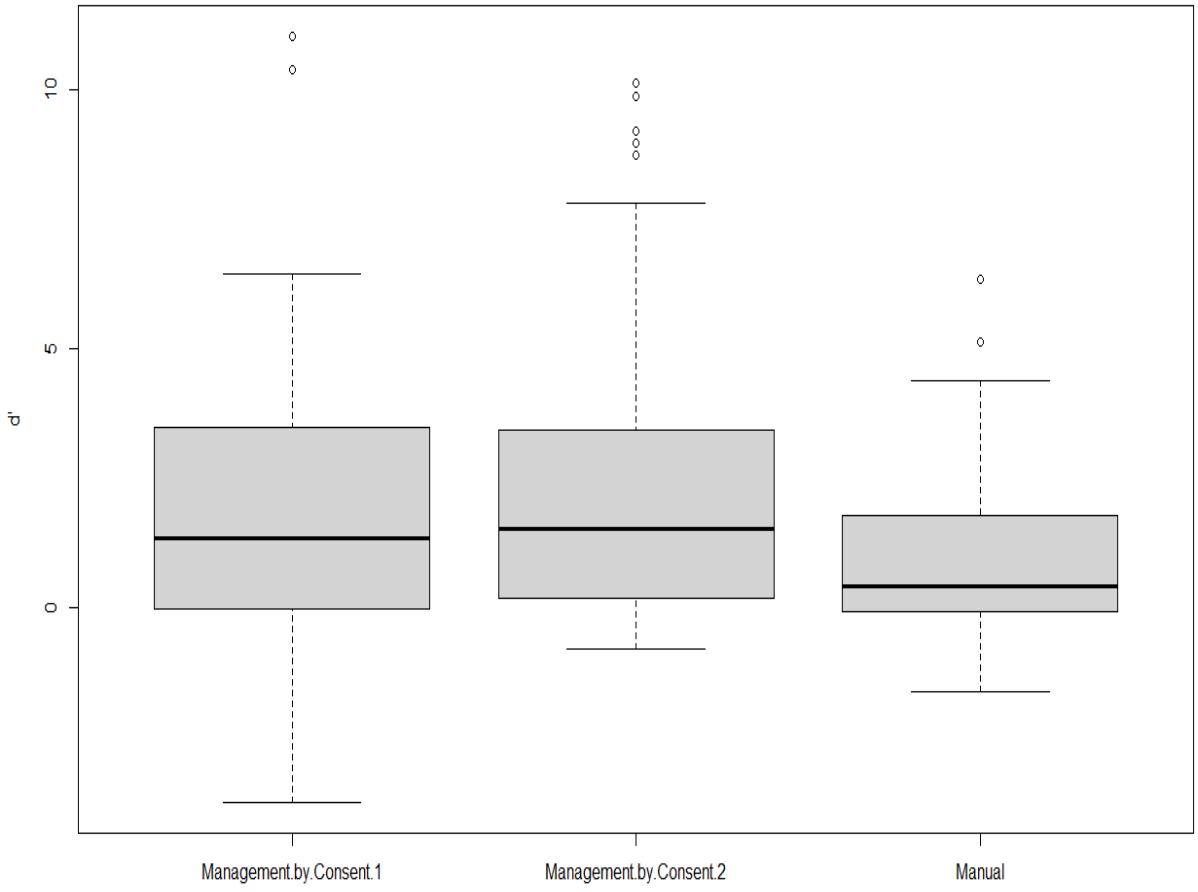

Figure 3: d' across Conditions

Table 2: Summary of Beacon Activation

Management-by-Consent1 (Polygon display of computer advice only)	146.08 (± 140.19)
Management-by-Consent2 (Polygon display plus air picture)	109.78 (± 107.67)
Manual Control (air picture only)	127.67 (± 123.06)

There was a significant main effect of condition on beacon activation [$X^2(2) = 12.98$, $p<0.005$]. Post-hoc pairwise comparison, using Wilcoxon, showed that Management by Consent 1 had significantly higher activation than Management by Consent 2 ($z=2.8$, $p<0.005$, power = 0.35) and Manual ($z=2.01$, $p<0.05$, power = 0.24). This is illustrated by table 2.

Table 3: Summary of Time per Target

Management-by-Consent1 (Polygon display of computer advice only)	0.295 (± 0.21)
Management-by-Consent2 (Polygon display plus air picture)	0.242 (± 0.24)
Manual Control (air picture only)	0.209 (± 0.24)

Table 3 indicates differences between conditions in terms of time-per-target. There is a significant main effect of condition on time per target ($X^2(2)=12.94$, $p<0.005$], and significant pairwise differences between all three conditions (table 4).

Table 4: Summary of Time per Target [Mean (sd)] *p<0.05, **p<0.001

	Auto (A) Management-by-Consent (computer advice)	Auto (A) Management-by-Consent (computer advice)	Auto (A) Management-by-Consent (computer advice)
Auto (A) Management-by-Consent (computer advice)	-	z = 2.15*	z = 3.46*
Both (B) Management-by-Consent (computer advice_air picture)		-	z = 2.03*
Self (S) Manual Control			-

Summary of Results

Participants were more conservative in their response when using the polygon display provided by Management by Consent 1. False Alarms in responding to the region around a protected house (FA_region) was lower for all conditions, which suggests that participants were sensitive to the collateral damage that could arise from an attack over the houses. However, false alarms in terms of attacks on non-hostile drones (FA_target) was significantly higher for the Manual condition, which also had more hits in total. This suggests in the Manual condition, participants were somewhat 'trigger' happy, in that UAS within range would more likely to be targeted. If this was the case, then one might expect the beacon activation measure to be highest for the Manual condition. However, the highest level of beacon activation was in the Management-by-Consent 1 condition. A possible explanation for this was that beacon activation provided a means of maintaining Situation Awareness through engagement in the game, in the absence of visual display of drone movement. If this was the case, then one might expect the time between destroying drones to be lower for the Management-by-Consent 1. In fact, we find that time per target is the *longest* for this condition. This suggests that, while beacon activation is higher for the Management-by-Consent 1 condition, this does not translate into more 'hits'. Rather, what seems to be happening is that the Management-by-Consent 1 condition involved increased beacon on / off activity (maintaining engagement), whereas beacon activation in the Manual Control condition resulted in more 'hits' (and false alarms) as participants responded to targets of opportunity.

Discussion

An initial observation is that participants in the Management-by-Consent 1 condition, i.e., when all they saw was the radar chart indicating the recommended response, still violated the rules. That is, UAS were still destroyed in the protected region. However, this was much lower than in the other conditions, and a plausible explanation might be that participants were simply not quick enough in deactivating the beacons (although the Management-by-Consent 1 condition resulted in significantly more beacon on / off commands). In this case, one might assume that fully automating the process could have minimized FA_ region to zero. However, the movement of the UAS was such that there were not always 'clear' opportunities to attack hostile targets and so there was often a dilemma in terms of managing collateral damage. Note, that providing a visual display of the moving drones (in the 'Management-by-Consent 2 and Manual Control conditions) has the effect of *increasing* the likelihood of hitting a UAS in a protected region, i.e., when participants were able to see the UAS and the protected regions, they seemed more likely to accept collateral damage. On the other hand, non-hostile UAS were still destroyed (FA_target). In this case, full automation might not have solved the problem because of the dynamic nature of the game; if you wait until only a hostile UAS is available then this could mean that targets of opportunity would be missed.

In future warfighting scenarios, attacks using multiple UAS could use dummy 'neutral' UASs to confuse the defence system. From this perspective, targeting hostile UAS when there are non-hostiles in the vicinity might be the only course of action to take; the challenge is to ensure that the damage to the non-hostile UAS is kept to a minimum. Note here, that our data reflect *all* 'hits' and that we do not record whether or not the UASs were flying with zero power; in a previous study, drones were destroyed and removed from the screen (Baber and Vance, 2020). However, it was felt that removing drones meant that the task become easier as the game progressed, which meant that there was little need of providing automated support). In this case, the strategy would be to offset damage, as reflected by the d' score. Here it is apparent that the Manual condition resulted in

significantly worse performance. This raises the question concerning the difference between awareness of the situation (e.g., the movement of the drones, the status of the beacons, the location of protected assets etc.) and the awareness of decisions (i.e., when to activate the beacon).

The study demonstrates that participants were able to monitor and respond to 15 UAS, as indicated by Ruff et al. (2002). The Management by Consent (in which participants were advised by automation but required to make a response) led to better compliance with the rules of engagement (in terms of minimising collateral damage) but lower 'hits'. In this case, the lack of an air picture (in Management by Consent 1 which only presented the polygon display) hampered participants ability to respond quickly. Indeed, the beacon activation data suggest that this resulted in a suboptimal strategy through which participants attempted to 'probe' the environment in order to check the movement of drones. When the air picture was presented, beacon activation was lower. We assume that this reflects improved Situation Awareness. However, we note that provision of the air picture alone resulted in a different form of suboptimal performance, in which sensitivity was significantly lower. In this respect, providing the decision support (in the form of the polygon display) helps to limit tendency for false alarms.

References

Baber, C., Morin, C., Parekh, M., Cahillane, M. and Houghton, R.J. 2011, Multimodal Control of Sensors on Multiple Simulated Unmanned Vehicles, *Ergonomics, 54,* 792-805

Baber, C. and Vance, C., 2020, Supporting decision making in a simulated air defence activity**,** In R Charles & D Golightly (Eds) *Contemporary Ergonomics & Human Factors 2020*

Botzer, A., Meyer, J., Borowsky, A., Gdalyhau, I. and Shalom, Y.B., 2015, Effects of cues on target search behavior, *Journal of Experimental Psychology: Applied, 21,* 73-88.

Chancey, E. T., Bliss, J. P., Yamani, Y., & Handley, H. A. H., 2017, Trust and the Compliance-Reliance Paradigm: The Effects of Risk, Error Bias, and Reliability on Trust and Dependence. *Human Factors, 59*(3), 333–345.

Cummings, M.L., Bruni, S., Mercier, S. & Mitchell, P.F. 2007, Automation Architectures for Single Operator, Multiple UAV Command and Control, *The International C2 Journal, 1,* 1-24

Liu, D., Wasson, R. & Vincenzi, D.A., 2009, Effects Of System Automation Management Strategies and Multi-Mission Operator-To-Vehicle Ratio on Operator Performance in UAV Systems, *Journal of Intelligent Robotics Systems, 54,* 795-810

Parasuraman, R., Sheridan, T.B. and Wickens, C.D., 2000, A model for types and levels of human interaction with automation, *IEEE Trans Sys, Man, Cyber – part A Systems and Humans, 30,* 286-297

Porat, T., Oron-Gilad, R., Rottem-Hovev, M. and Silbiger, J., 2016, Supervising and controlling unmanned systems: a multi-phase study with subject matter experts, *Frontiers in Psychology, 7,* 56.

Ruff, H.A., Narayanan, S. and 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,* 335-351.