Human Augmentation of UAV Cyber-Attack Detection

Haibei Zhu[1], Mahmoud Elfar[1], Miroslav Pajic[1], Ziyao Wang[2] and M. L. Cummings[2]

[1] Department of Electrical & Computer Engineering, Duke University, Durham NC 27708
[2] Department of Mechanical Engineering & Materials Science, Duke University, Durham NC 27708
{haibei.zhu, mahmoud.elfar, miroslav.pajic, ziyao.wang, mary.cummings}@duke.edu

Abstract. Unmanned aerial vehicles (UAVs) have extensive applications in both civilian and military applications. Nevertheless, the continued development of UAVs has been accompanied by security concerns. UAV navigation systems are potentially vulnerable to malicious attacks that target their Global Positioning System (GPS). Thus, efficient GPS hacking detection with high success rate is paramount. Significant effort has been put into developing autonomous hacking detection techniques. However, little research has considered how a human operator can contribute to the security of such systems. In this paper, we propose a human-autonomy collaborative approach for a single operator of multiple-UAV supervisory control systems, where human geo-location is used to help detect possible UAV cyber-attacks. An experiment was designed and conducted using the RESCHU-SA experiment platform to evaluate this approach. The primary results show that 65% of all experiment sessions reached over 80% success rate in UAV hacking detection, while only 17% of participants lost one or more UAVs because of incorrect hacking detections. These results suggest that such an approach could help achieve better security guarantees for human-in-the-loop autonomous UAV systems that are prone to cyber-attacks.

Keywords: Unmanned Aerial Vehicles, Cyber-Attack Detection, Human Geo-location.

1 Introduction

Unmanned aerial vehicles (UAVs) have significantly increasing commercial market and extensive applications in both civilian and military realms [1]. Many of these UAVs rely on the Global Positioning System (GPS) for navigation, however, this reliance leaves UAVs vulnerable to malicious attacks targeting GPS signals. One common attack is GPS spoofing, in which attackers deceive GPS receivers to override the navigation systems and redirect UAVs to unexpected destinations [2] [3]. A well-known such incident garnered public attention in 2011 when, a US RQ-170 Sentinel UAV was captured by Iranian forces using GPS spoofing attacks [4]. Thus, detecting GPS spoofing attacks with a high success rate is important for UAV control systems.

We propose a human-autonomy collaborative approach of human geo-location in that humans can aid in the detection of possible GPS spoofing attacks on UAVs. This

approach was evaluated via an experiment, which was designed and conducted using the Research Environment for Supervisory Control of Heterogeneous Unmanned Vehicles (RESCHU) platform. Experiment sessions simulated human supervisory multi-UAV control scenarios with potential UAV GPS spoofing attacks. In this paper, we focus on answering the following questions based on the experiment results: 1) Can human operators successfully identify UAV GPS spoofing attacks? 2) What factors affect human operator general operation? 3) Would hacking detections affect the performance of operators' primary tasks? 4) What types of landmarks used in human geo-location affect operator decisions to hacking detections?

2 Background

A common UAV control scheme is human supervisory control, in which a human operator monitors the multi-UAV system, intermittently navigating UAVs, and conducting other higher-level tasks [5]. The architecture of human supervisory UAV control is shown in Figure 1. Human supervisory UAV control can be introduced with various level of automation. In this study, we assume that human operators are responsible for higher-level decision, and autonomous systems are in charge of lower-level UAV control and navigation operations [6].

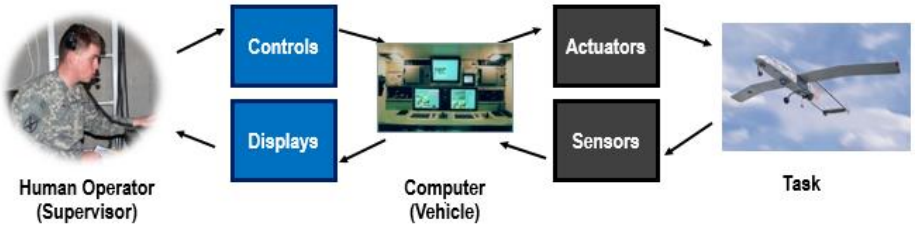

Fig. 1. Human supervisory UAV control architecture.

2.1 GPS Spoofing Detection

UAVs typically rely on an embedded navigation system known as GPS, which provides accurate position, velocity and time information for GPS receivers in most areas on Earth. GPS receivers calculate precise latitude, longitude and height with speed information based on the received satellite signals. Furthermore, GPS receivers can report their locations to UAV control interface to provide location views for operators. However, GPS receivers are vulnerable to GPS spoofing attacks, in which GPS spoofers generate counterfeit signals to attack GPS receivers by manipulating the target position, velocity and time information [2] [3].

Many researchers have presented autonomous GPS spoofing detection methods [7] [8] [9] [10] [11] [12], however, false alarms and detection mistakes still exist while

applying autonomous detection techniques [13] [14]. Thus, supplementary detection methods are needed.

In the common design of military UAVs, a UAV is usually equipped with both a GPS navigation system and a payload camera, whose signal is independent of the UAV GPS signal [15]. Thus, the UAV payload camera view could be used as an independent reference for detection of GPS spoofing (i.e., navigation based) attacks, which is further explored in the remainder of this paper.

2.2 Human Visual Task

In order to utilize a UAV payload camera to detect UAV GPS attacks, interpreting the UAV real-time location through the camera view and comparing this to a certain landmark or position estimate from a map could be the central mechanism for making such an assessment.

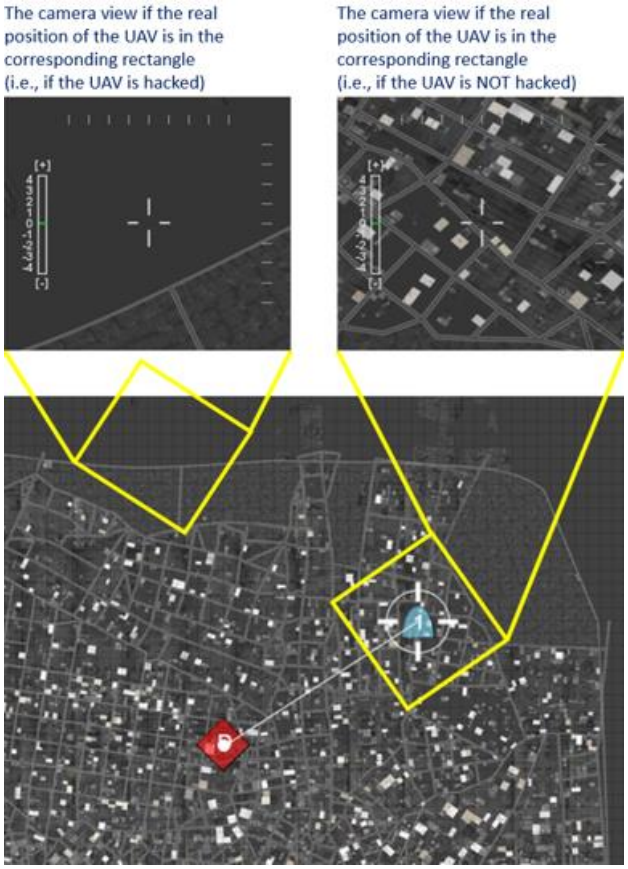

Fig. 2. An example of GPS reported location on the map.

While autonomous localization techniques may have limited performance [16] [17], human vision has advantages in such complex search and surveillance tasks. The process of human vision obtaining information from objects can be divided into two stages. The first stage is the preattentive stage, in which human observers can gather basic information about the target even before the observer become conscious of it [18]. Thus, human vision can process target information relatively fast in complex environment. Human observers also tend to choose areas that maximize information of the target in salience-driven visual search strategy [19], which means human vision has effective strategies to obtain target information. In addition, the direction discrimination threshold of human vision has a low average of 1.8 degrees [20], which means human vision can detect relatively small changes in orientation. Based on these visual advantages, a human operator can potentially aid in UAV localization and thus detect potential UAV GPS spoofing attacks.

Based on the assumption that UAV cameras can show the true surrounding scene of UAVs, we propose that human operators can act as supplementary sensors and assist autonomous system to detect UAV hacking attacks through comparative geo-location between the camera and map position estimates. In human geo-location, the operator can compare the non-tempered video feed coming from the UAV to the potentially falsified GPS location; this allows the user to detect inconsistencies between these two sensing feeds (i.e., whether the feed and the reported locations match). If the operator thinks the location interpreted from camera view does not match the location shown on the map, then the UAV is most likely hacked via GPS spoofing.

An example of applying human geo-location in UAV hacking detection is shown in Figure 2. If the UAV is hacked, the operator will observe a location other than the GPS reported UAV location through the camera, like shown in the upper left camera view in Figure 2. When a GPS spoofing attack is discovered, the operator can prevent a hacking event by overriding its physical control.

3 Experiment

An experiment was designed utilizing a modified version of the RESCHU experiment platform [21], known as Security-Aware RESCHU (RESCHU-SA) [22]. RESCHU-SA is a Java-based single operator with multi-UAV supervisory control simulation platform, which provides the capability to design multi-tasking scenarios that include both navigational and imagery search tasks. Moreover, the platform allows for simulating GPS spoofing attacks, in which hacked UAVs deviate from their originally assigned path and target to other unexpected destinations, along with warning notifications that simulate autonomous GPS spoofing detection systems.

3.1 Experiment Platform Interface

The interface of the RESCHU-SA platform is shown in Figure 3. The interface features five main components: the payload camera view, message box, control panel, mission timeline and map area.

- The camera view displays the video stream from the payload camera of the selected UAV. The primary purpose of this view is to conduct real-time image analysis tasks. In this study, it can also be used to determine the actual location of UAVs by locating landmarks.
- The message box displays events that occur during the simulation such as UAV arrival at a target. It also allows operators to communicate the results for the imagery analysis tasks to a "supervisor" that is, in actuality, a bot.
- The control panel provides the UAV damage level, which is caused by UAVs intersecting with hazard areas, as well as instructions for imagery analysis tasks and vehicle updates.
- The timeline shows the estimated remaining time of all UAV arrivals at waypoints and assigned targets.
- The map displays the area of surveillance with real-time locations of all UAVs, hazard areas and targets. For this experiment, the map was created using CityEngine from ArcGIS, a modeling software package that is used for urban planning and architecture design.

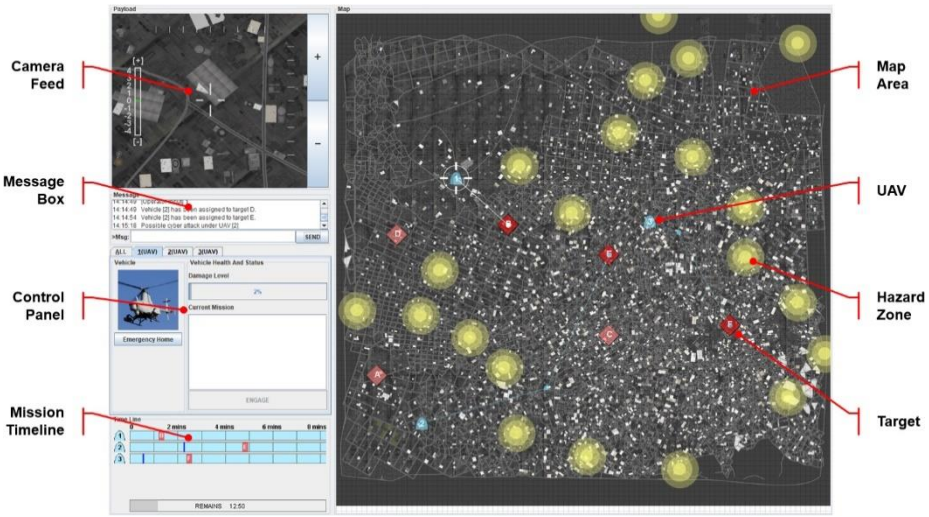

Fig. 3. RESCHU-SA operator interface.

3.2 Experiment Scenarios

The primary objectives for operators in RESCHU-SA are to control multiple UAVs to: 1) perform reconnaissance imagery tasks of counting road intersections when UAVs reach assigned targets, 2) ensure UAVs do not encounter hazard areas, and 3) determine whether UAVs are under GPS spoofing attacks.

For this experiment, GPS spoofing attack events followed a pre-defined schedule, unknown to the participants. When triggered internally, the hacked UAV changed its heading by a random angle within 30 to 60, or 300 to 330, degrees, which was larger

than the human direction discrimination threshold. A hacking notification appeared 10 to 20 seconds after the attack event, simulating an external agency detection of a possible GPS spoofing attack. However, as in real systems, the notification could be a false alarm. In fact, about half of all notifications in this experiment were false alarms in the pre-defined schedule of each test session. Although change blindness and vigilance declination [23] [24] are known problems for humans in such settings, they were not examined in this study, given there was a clear signaling of a potential hacking event. Thus, a detection failure on the part of the human is not considered in this experiment.

Once the operator received notification from the system that a certain UAV was under possible cyber-attack, the operator could then investigate the UAV by checking the UAV's camera view and matching it against the position of the UAV on the map. The operator was expected to make a decision before the hacked UAV either exceeded the map boundary or the experiment ended. If the operator decided the UAV was hacked, the operator could override the hacked UAV and send it home.

When UAVs that were not hacked reached a target, the operator engaged in an imagery task of counting the road intersections from the UAV's camera view at a prespecified zoom level. This side task represents the primary purpose of such a mission, which is typically information gathering. While engaging in a counting task, the operator was required to enter an answer before the counting task was finished. The counting task allowed us to assess participants' performance based on the number of attempted tasks and the task correctness percentage.

The path planner for the UAVs was intentionally suboptimal, i.e. the planner did not necessarily pick the most efficient assignment of UAVs to targets. In addition, UAVs would possibly encounter hazard areas that appeared and disappeared randomly. The suboptimal planner and the dynamic nature of hazard areas allowed experimenters to assess how much spare attention participants could devote to optimize the navigation and target assignment.

3.3 Experiment Participants

Thirty-six participants took part in this experiment, including 22 males and 14 females. Age ranged from 19 to 34 years with an average of 25.2 and a standard deviation of 3.8 years. Among the participants, 18 had little video game experience, 6 participants had monthly gaming experience, 5 participants played video game several times a week, another 5 participants had weekly gaming experience, and only 2 participants had daily gaming experience.

3.4 Experiment Procedure

The experiment procedure consisted of four main sections. The first section was a self-paced tutorial session, during which participants went over the tutorial slides, and the experimenter answered questions that the participants might have had. The second section was a practice session to allow participants to get more familiar with the user

interface. In the first half of the practice session, participants were shown how to operate UAVs and complete all major tasks.

In the second half, participants controlled all UAVs and accomplished different tasks by themselves. The practice session lasted 18 minutes, which was the same as a single experiment session. The third section included the test sessions with two scenarios of different task loads, which were counterbalanced in terms of order of presentation. The fourth section was the debriefing session, in which the experimenter asked the participant several questions related to participants' performance and strategies for navigating UAVs and detecting hacking events.

Given that many related studies on the RESCHU platform [21] [25] [26] showed a significant impact of task load on system performance, task load was a primary factor in this experiment looking at hacking detection. It should be noted that high task load does not necessarily represent high operator mental workload, since operator mental workload is an individually subjective interpretation of an objective task load.

Thus, for a high task load scenario, operators controlled 6 UAVs with 9 different targets and 9 hacking events, and in each low task load scenario, operators controlled 3 UAVs with 6 different targets and 6 hacking events. In both scenarios, the number of hazard areas, which generated and disappeared randomly, was constantly twenty-one. Each test scenario lasted 18 minutes, and each participant completed both high and low sessions. Each participant's performance scores were calculated based on the total vehicle damage, the correct percentage of imagery counting tasks, and the correct percentage of hacking identifications.

4 Results

4.1 Performance statistical results

We used a multivariate repeated-measures ANOVA model and Pearson correlation with a significance level of 0.05 to analyze data. In data analysis, independent variables included task load, which task load was experienced first, gender and video game experience as a covariate. Task load (low and high) was a within factor variable. Dependent variables included percentage of correct hacking detections, the aggregated damage sustained by vehicles over a test session, and the overall correct percentage intersection counts per test session. These variables represent the primary objectives of performing the counting tasks, keeping vehicles out of the damage areas, and successfully detecting hacking events.

Table 1. The confusion matrix of hacking detection decisions in different notifications.

	Real hacking notification	False alarm notification
Consider UAV hacked	224	40
Consider UAV not hacked	63	207

An important question was whether human operators could successfully detect the UAV hacking events. A successful detection was indicated by a correct decision for a specific hacking event, including overriding the UAV and sending it home if the UAV was hacked, or recognizing the notification was a false alarm.

Each high task load experiment session included 9 hacking events, and each low task load session included 6 hacking events. Among all hacking events in both test sessions for each participant, 7 (4 in high task load and 3 in low task load) were predefined as false alarms, which meant the threshold for incorrect hacking notifications was 47%. As shown in Table 1, out of all real hacking notifications across all participants, the overall success rate was 78%, and for the false alarms, the success rate was 84%. In other word, the type one error (false positive, operators considered UAV not hacked with real hacking notification) was 22%, which was slightly higher than the type two error (false negative, operators considered UAV hacked with false alarm notification) of 16%. Thus, operators were slightly better at detecting false alarms than identifying real hacking notifications.

When looking at each individual's performance per test session, even though they had to multitask in RESCHU-SA in managing multiple vehicles and detecting potential hacking events, results showed that 23 out of total 72 experiment sessions (32%) resulted in 100% of successful hack identifications in a single test session, with another 24 (33%) above 80% successful attack identification. Thus, 65% of total experiment sessions exhibited 80% correct hacking detection or better without having any prior formal training. In terms of incorrect hacking identifications, 12 (17%) participants lost one or more UAVs, meaning that these UAVs were successfully hacked and could not be further controlled.

Additionally, those factors that affected human operators' performance were studied. For the three performance scores of vehicle damage, the correct percentage intersection counts, and correct percentage of hacking events, the only variable affected by task load was vehicle damage ($(F(1,31)=32.93)$, $p<0.001$). In the high task load scenario, the average UAV damage was 31.4, which was much higher than 9.6 in the low task load scenario. Participants with less workload suffered less damage as they had more time to optimize their paths and avoid hostile areas.

One result showed an interesting significant negative correlation between the time expended in hacking detections and correct hacking detections (Pearson=-0.375, $p=0.001$), which meant that participants who took longer to detect the hacking events had a lower percentage of correct hacking identifications. This suggests that early detection was better from the operator standpoint, which is at odds with those who would argue that longer detection times should yield more correct identifications.

Gender was examined because of the potential difference in self-assessment in cognitive tasks between different genders [27]. However, gender did not affect the participants' general performance. Another covariate, the video game experience, did have a significant effect on participants' correct hacking detections ($F(1,31) = 4.652$ $p = 0.039$). This means that the more video game experience, the higher the chance of a correct hacking detection. Not surprisingly, seven participants who lost UAVs had no video game experience, and the other 5 who lost UAVs ranged from some to moder-

ate gaming experience. Participants with daily gaming experience did not lose any UAVs and were 100% correct in hacking identification.

Another result showed that participants' task inputs were effective in that the more time they navigated the UAVs, the less time UAVs intersected with hostile areas (Pearson=-0.345, p=0.003). This result suggests that improved path planning could reduce operators' workload and free their cognitive resources to attend to other tasks.

We also investigated whether hacking detection affected the performance of operators' primary tasks of counting road intersections. The results showed that the correctness of imagery counting tasks was not affected by either the correctness of hacking detections (Pearson=-0.022, p=0.854) or the time expended in hacking detections (Pearson=0.024, p=0.841). However, time expended in the imagery task was negatively correlated with the percentage of correct hacking detection (Pearson=-0.275, p=0.019). This result was expected as participants who spent more time in counting tasks were less likely to detect hacking events.

In addition, an interesting observation is that the first experiment scenario affected participants' abilities to correctly finish their primary task of counting the intersections at each target ($F(1,31)=5.324$, $p=0.028$). The participants who had the high task load scenario as the first experiment session tended to have higher correct intersection count percentages. This suggests a fatigue effect since these participants did worse on their second scenarios with low task load, which should have been easier.

4.2 Map analysis for hacking detection

While using human geo-location in UAV hacking detections, operators will compare the non-tempered UAV camera video feed to the potentially falsified GPS location to detect inconsistencies between the UAV video feed and UAV GPS location. After receiving a hacking notification, operators can purposely navigate the notified UAV to some specific areas that can potentially provide more inconsistencies to increase the confidence of making a correct decision to a hacking event. Thus, analyzing the map usage in hacking detections will benefit the future design of autonomous decision supporting tool for hacking identification.

The resulting heat map, which represents the frequency distribution of areas of participant interest during hacking detections, is shown in Fig 4. Different colors represent varying frequency of operations, including adding waypoints and switching targets for UAVs, on a specific point. The warmer the color, the more participants interacted with a specific point, for example, red represents 5 or 6 operations. The heat map shows that the lower left quadrant is the most popular region, however, some regions, like the middle right of the map, have few operations. Understanding that the density of targets on the lower left quadrant of the map is slightly higher than other regions, this quadrant is more attractive to operators since operators can navigate UAVs between targets to get engaged to more imagery tasks in a shorter time range. Thus, more operations occurred on the lower left quadrant. In addition, red areas on this quadrant indicate the existence of some interesting landmarks that operators tend to investigate during hacking detections.

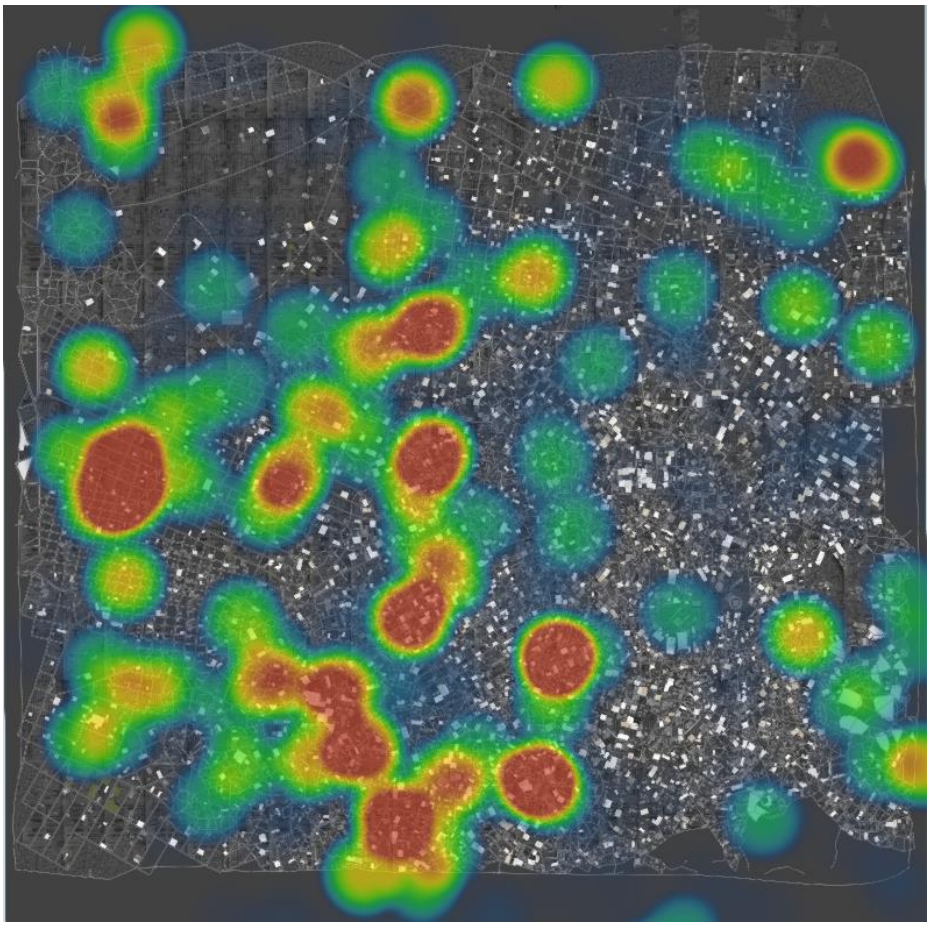

Fig. 4. The heat map of reference points in UAV hacking detection.

Table 2. The frequency of different types of landmarks used in hacking detections.

	Special road patterns	Geographic feature transition	Special buildings
Occurrence frequency	380	152	109
Occurrence percentage	59.3%	23.7%	17.0%

Landmarks used in hacking detections are classified into three categories, including special road patterns, geographic feature transition, and special buildings, like shown in Table 2. Using these different landmarks in hacking detections, operators can investigate the moving orientation of a certain UAV or the relative motion between a UAV and a specific landmark to investigate whether a UAV is potentially hacked. Shown in Table 2, special road patterns were the most frequently used landmarks in hacking detection with an occurrence percentage of 59.3%.

Geographic feature transitions are defined as the transition between land and sea areas, on which operators can clearly observe the sudden change of geographic patterns. Special buildings are defined as distinctive shapes with contrastive colors that are used to represent a single building or a group of buildings on the map. As the percentage of total special road patterns and special buildings are approximate the same, special road patterns are more attractive to operators. Future work will determine why exactly people prefer these over other options, but one hypothesis is that these are easier to see than the buildings, and do not take as long to investigate as the sea/land transitions.

Table 3. The frequency of different landmarks used in different detection decisions.

	Real hacking notification		False alarm notification	
	Consider UAV hacked	Consider UAV not hacked	Consider UAV hacked	Consider UAV not hacked
Special road patterns	168	25	28	159
Geographic feature transitions	69	17	14	52
Special buildings	42	13	8	46

The frequency of different landmarks used in different detection decisions were examined. In correct hacking detections with both real hacking and false alarm notifications, the percentage of operations based on special road patterns is slightly over 60%, which is higher than the percentage in incorrect hacking detection with real hacking notification (45.5%) and false alarm notification (56.0%). Another interesting fact is that special road patterns lead to the highest success rate of 86.1% in hacking detections, while geographic feature transition lead to 79.6% and special buildings lead to 80.7%. These results provide insight for how a future advanced map-based hacking detection support tool for human operators could be designed.

5 Discussion

The experiment results provide insight into our initial questions with implications for future studies. In this study, we analyzed if a human operator could serve as a supplementary sensor in supervisory UAV control systems by successfully detecting spoofing attacks. Experiment results supported this hypothesis in that 65% of total experiment sessions reached over 80% hacking detection correctness. This result was achieved with no dedicated training and so greater emphasis on optimal search strategies would likely yield even better results.

The experiment results also clearly indicate that some factors affected operators' performance and operations. For example, higher task load tended to cause more UAV damage. This result was supported by a previous study that higher mental work-

load increased operator attention switching delays [21]. In high task load scenarios, operators tended to experience higher mental workload, which slowed down their attention switches and causing more damage. This could be mitigated in future studies with more optimal path planning as well as better target allocation.

Understanding that the operator's video game experience significantly affected the success rate in hacking detections, future personnel selection strategies for supervisory control systems with human visual tasks could focus more on the experience in similar applications or more training. This fact also raises interesting future research questions, including how video game experience may affect human search strategies and how different types of video games may affect human operators' performance in hacking detection? Also, the result of the negative correlation between the time expended and the success rate in hacking detection provides implications for future studies of increasing hacking detection correctness by guiding better search strategies and earlier detections. However, a fatigue effect was potentially exhibited after just one 18-minute scenario, which raises the question of how sustainable such task load levels are over time?

The map analysis shows the heat map of participants preferences for hacking detection. Although the usage percentages of different landmarks in different hacking detection decisions are similar, there was a clear preference for unusual road intersections. These results provide some insights on a more efficient way to utilize different landmarks and raise future research topics of investigating potential different operator hacking detection strategies.

Lastly, all these results establish a baseline of performance of applying human geo-location in UAV hacking detection. Future studies, enabled by an empirical model of security-aware human-autonomy interaction will focus on how higher-level automation or advanced decision support tools could be utilized to assist human operators to improve the success rate of hacking identifications.

6 Conclusion

Navigational GPS systems used in UAVs can be prone to malicious cyber-attacks, especially GPS spoofing attacks. In this study, we have shown that a human operator can assist autonomous systems in hacking detection using human geo-location comparison between maps and downward-facing camera views, even without extensive training. Moreover, we found that an individual factor, video game experience, and the time expended in hacking detection and UAV navigation, affected operators' hacking detection performance. The results from this study indicate that human geo-location is a potentially promising approach for hacking detection, which could be improved by future efforts in improving operator decision support.

Acknowledgements. This work was supported in part by the CNS-1505701 grant. This material is also based on research sponsored by the ONR under agreements number N00014-17-1-2012 and N00014-17-1-2504. We gratefully acknowledge the efforts of those who have assisted in developing the RESCHU-SA experiment plat-

form including Duke undergraduate students, Jeffrey Wubbenhorst, Adithya Raghunathan and Yi Yan Tay.

References

1. Pajares, G. (2015). Overview and current status of remote sensing applications based on unmanned aerial vehicles (UAVs). Photogrammetric Engineering & Remote Sensing, 81(4), 281-329.
2. Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O'Hanlon, B. W., & Kintner Jr, P. M. (2008, September). Assessing the spoofing threat: Development of a portable GPS civilian spoofer. In Proceedings of the ION GNSS international technical meeting of the satellite division (Vol. 55, p. 56).
3. Shepard, D. P., Humphreys, T. E., & Fansler, A. A. (2012). Evaluation of the vulnerability of phasor measurement units to GPS spoofing attacks. International Journal of Critical Infrastructure Protection, 5(3), 146-153.
4. Shane, S., & Sanger, D. E. (2011). Drone crash in Iran reveals secret US surveillance effort. The New York Times, 7.
5. Sheridan, T. B. (1992). Telerobotics, automation, and human supervisory control. MIT press.
6. Cummings, M. L., Bruni, S., Mercier, S., & Mitchell, P. J. (2007). Automation architecture for single operator, multiple UAV command and control. Massachusetts Inst Of Tech Cambridge.
7. Kerns, A. J., Shepard, D. P., Bhatti, J. A., & Humphreys, T. E. (2014). Unmanned aircraft capture and control via GPS spoofing. Journal of Field Robotics, 31(4), 617-636.
8. Wesson, K. D., Shepard, D. P., Bhatti, J. A., & Humphreys, T. E. (2011, September). An evaluation of the vestigial signal defense for civil GPS anti-spoofing. In Proceedings of the ION GNSS Meeting.
9. Broumandan, A., Jafarnia-Jahromi, A., Dehghanian, V., Nielsen, J., & Lachapelle, G. (2012, April). GNSS spoofing detection in handheld receivers based on signal spatial correlation. In Position Location and Navigation Symposium (PLANS), 2012 IEEE/ION (pp. 479-487). IEEE.
10. Wesson, K., Rothlisberger, M., & Humphreys, T. (2012). Practical cryptographic civil GPS signal authentication. Navigation, 59(3), 177-193.
11. Psiaki, M. L., O'Hanlon, B. W., Bhatti, J. A., Shepard, D. P., & Humphreys, T. E. (2013). GPS spoofing detection via dual-receiver correlation of military signals. IEEE Transactions on Aerospace and Electronic Systems, 49(4), 2250-2267.
12. Pajic, M., Weimer, J., Bezzo, N., Sokolsky, O., Pappas, G. J., & Lee, I. (2017). Design and Implementation of Attack-Resilient Cyberphysical Systems: With a Focus on Attack-Resilient State Estimators. IEEE Control Systems, 37(2), 66-81.
13. Wesson, K. D., Evans, B. L., & Humphreys, T. E. (2013, December). A combined symmetric difference and power monitoring GNSS anti-spoofing technique. In Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE (pp. 217-220). IEEE.
14. Humphreys, T. E. (2013). Detection strategy for cryptographic GNSS anti-spoofing. IEEE Transactions on Aerospace and Electronic Systems, 49(2), 1073-1090.
15. Sun, J., Li, B., Jiang, Y., & Wen, C. Y. (2016). A Camera-based target detection and positioning UAV system for search and rescue (SAR) purposes. Sensors, 16(11), 1778.
16. Radke, R. J., Andra, S., Al-Kofahi, O., & Roysam, B. (2005). Image change detection algorithms: a systematic survey. IEEE transactions on image processing, 14(3), 294-307.

17. Blacknell, D., & Griffiths, H. (2013). Radar automatic target recognition (ATR) and non-cooperative target recognition (NCTR). The Institution of Engineering and Technology.
18. Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. Cognitive psychology, 12(1), 97-136.
19. Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. IEEE Transactions on pattern analysis and machine intelligence, 20(11), 1254-1259.
20. De Bruyn, B., & Orban, G. A. (1988). Human velocity and direction discrimination measured with random dot patterns. Vision research, 28(12), 1323-1335.
21. Donmez, B., Nehme, C., & Cummings, M. L. (2010). Modeling workload impact in multiple unmanned vehicle supervisory control. IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 40(6), 1180-1190.
22. Elfar, M., Zhu, H., Raghunathan, A., Tay, Y. Y., Wubbenhorst, J., Cummings, M. L., & Pajic, M. (2017, April). Platform for security-aware design of human-on-the-loop cyber-physical systems. In Proceedings of the 8th International Conference on Cyber-Physical Systems (pp. 93-93). ACM.
23. Simons, D. J., & Ambinder, M. S. (2005). Change blindness: Theory and consequences. Current directions in psychological science, 14(1), 44-48.
24. Temple, J. G., Warm, J. S., Dember, W. N., Jones, K. S., LaGrange, C. M., & Matthews, G. (2000). The effects of signal salience and caffeine on performance, workload, and stress in an abbreviated vigilance task. Human factors, 42(2), 183-194.
25. Boussemart, Y., & Cummings, M. L. (2011). Predictive models of human supervisory control behavioral patterns using hidden semi-Markov models. Engineering Applications of Artificial Intelligence, 24(7), 1252-1262.
26. Boussemart, Y., Cummings, M. L., Las Fargeas, J. C., & Roy, N. (2011). Supervised vs. Unsupervised Learning for Operator State Modeling in Unmanned Vehicle Settings. JACIC, 8(3), 71-85.
27. Pallier, G. (2003). Gender differences in the self-assessment of accuracy on cognitive tasks. Sex Roles, 48(5), 265-276.