Finnish Defence Research Agency
Research Bulletin 2026–5

24.3.2026

Human-Swarm Interaction

Salla Skön, Lauri Oksama
Human Performance Division

The introduction of unmanned, remotely operated robots has revolutionised modern warfare in the recent years. This review explores factors affecting the performance of swarm operators. Situation awareness, cognitive load and trust in automation are identified as the most prominent human factors issues in human-swarm interaction. Proper user interface design and important skills for swarm operators are also highlighted, along with focus points for future research.

Introduction

The introduction of unmanned, remotely operated robots has revolutionised modern warfare in the recent years. In June 2025, Ukraine carried out an unprecedented drone operation (Spider's Web) on four strategically important airfields in Russia (Bondar, 2025). Ukraine managed to destroy approximately ten and damage more than forty aircraft during the operation. According to estimates, this amounts to one third of Russia's strategic cruise missile delivery platforms (Bondar, 2025). Ukraine employed 117 drones in total, costing under two thousand dollars apiece. Therefore, with only moderate investments in equipment, Ukraine was able to cause Russia losses of roughly seven billion dollars (C. Walker, 2025). Ukraine has also reported that all personnel involved in the operation were able to return to Ukraine safely (Gozzi & Verify, 2025).

The operation described above showed that inexpensive unmanned systems can effectively destroy extremely valuable enemy assets. However, this operation did not use a proper drone swarm but separate drones each controlled by an operator (Gozzi & Verify, 2025). By combining several unmanned vehicles to a swarm, the number of operators needed to control the same number of vehicles is reduced. Consequently, using swarms might be one of the most cost-effective means to increase the firepower of a single soldier.

Research on swarms is still heavily nested in engineering sciences. Most studies focus on the technical development of swarm systems, ignoring the usability from the operator's standpoint (Hocraffer & Nam, 2017). Nevertheless, it is extremely important to investigate the limits of human information processing regarding human-swarm interaction – they will not change even though technology keeps on developing. The purpose of this report is to produce information on human performance when working with swarms, as well as to examine ways to bypass the restrictions defined by human information processing capacity in order to achieve efficient human-swarm interaction.

Methods

The material was collected using search words ("human-swarm interaction" OR "human-agent interaction" OR "human-multi-agent systems") AND ("cognitive load" OR "cognitive processing" OR "cognitive constraints" OR "human cognition") from the last 15 years. The databases used in this research were IEEE Explore, ACM Digital Library, APA PsycINFO, Science Direct, SpringerLink, ProQuest Military Database and EBSCO. The material was supplemented with relevant papers during the writing process.

What is a swarm?

Human interaction with single unmanned vehicles has been studied widely (see e.g., Nisser & Westin, 2006; Özyörük, 2020). Also, human interaction with multi-agent teams, which consist of several, separate agents, has drawn lots of attention (e.g., Chen & Barnes, 2014; Dixon et al., 2005). In traditional multi-agent-systems, performance is quickly impaired as the number of agents increases and the operator's information processing capacity is exceeded. One operator can control a maximum of four separate vehicles simultaneously without performance decaying excessively (Cummings & Mitchell, 2008). Swarming helps bypass this limitation by enabling the handling of several co-operating robots as a single, coherent entity (Kolling et al., 2015).

The word "swarm" is used quite loosely in scientific literature, usually in lieu of "multi-agent-system". The actual definition of a swarm is a large group of entities (e.g., drones) that move and act in synchrony. A swarm is not necessarily centrally controlled, and individual agents do not have complete knowledge of the swarm's goals or state. Local interactions between robots causes the emergent behavior typical to swarms, and it cannot be explicitly defined or planned (Bjurling, 2025). This swarm-characteristic behavior also results in a high fault tolerance (Crandall et al., 2017). In a swarm robots can navigate, make decisions and coordinate actions autonomously. Objects in the vicinity (including adjacent robots) influence the whole swarm's behavior (Kolling et al., 2015). Even though there is some level of autonomy, human operators still play a crucial part in managing swarms: steering the swarm toward strategic goals, communicating information to the swarm, and aiding the swarm in flexible and complex environments all belong to the operator's responsibilities, not to forget ethical issues (Crandall et al., 2017).

On the one hand, swarm systems have yielded remediation to some weaknesses in multi-agent-systems. On the other hand, novel human factors challenges have emerged due to the idiosyncrasies of swarms.

Human factors challenges

In military, there is constant pressure to increase the number of vehicles a single operator can control. However, the challenges associated with human performance are exacerbated with increasing complexity and size of the swarm. Challenges arise from the cognitive limitations of humans, technological restrictions, and suboptimal interfaces, among other things (Hocraffer & Nam, 2017). In addition to these issues, swarms considerably differ from autonomous agents or traditional robots. To understand how a swarm functions, the operator needs to comprehend system-level dynamics instead of focusing on the actions of individual robots (Kolling et al., 2015). Future systems must be effortlessly scalable to larger swarms without imposing unreasonable workload on the operator controlling the swarm.

HFACS (Human Factors Analysis and Classification System) model can be used to structually analyse human errors (e.g., error types and their causes) (Shappell & Wiegmann, 2000). Alharasees & Kale (2024) have studied HFACS in the context of unmanned vehicles. These same factors can be considered essential also when operating swarms. The most common human factors errors individuals make are related to decision-making (e.g., as a result of too much cognitive load), skills (e.g., inadequate training), perception (e.g., incorrect interpretation of data) and violations (e.g., intentionally disregarding AI protocols) (Alharasees & Kale, 2024).

Aforementioned error types can be summarised into three most essential human factors categories: situation awareness, cognitive load and trust in automation. This report focuses on research data relevant to them. There are various other error types (e.g., organisational, supervisory, environmental), but they are beyond the scope of this report.

Situation awareness

Situation awareness (SA) entails the operator maintaining real-time knowledge of the state of the swarm, mission and environment, understanding the meaning of this information, and utilising it in decision-making. The ability to perceive the state and movement of the swarm as well as to predict the swarm's future state with reasonable accuracy are especially highlighted when controlling semi-autonomous swarms (Kolling et al., 2015).

Situation awareness has a significant effect on decision-making and therefore on task performance. Considering remotely controlled robots, one of the most important factors impairing situation awareness is difficulties in communication between the operator and the swarm (Parush et al., 2011). Other factors reducing situation awareness are inappropriate user interfaces and ways to display information (Rojas et al., 2020). Sometimes the use of automation can also hinder situation awareness, even though it might help reduce cognitive load (Rojas et al., 2020).

Time perception is also essential regarding situation awareness. The operator needs to know when something task-relevant is about to happen and to be able to plan the optimal timing to execute actions (Endsley & Garland, 2000). Humans have the tendency to underestimate elapsed time when monitoring swarm movements (Elkin-Frankston et al., 2023). Bigger swarm size diminishes perceived time, whereas slower robot movement inflates it (Kaduk et al., 2023). Kaduk et al. suggest that future swarm systems could autonomically modify subjective time perception in order to ease the operators' work. This could be especially beneficial in, for example, critical scheduling tasks (Kaduk et al., 2023). The exact significance and uses of modifying subjective time perception is yet to be defined through research.

A widely used method for measuring SA is the SAGAT (Situation Awareness Global Assessment Technique) questionnaire, where a task simulation is interrupted at random points, and the operator answers questions measuring real-time situation awareness (Endsley, 1988). Situation awareness and cognitive load are strongly linked to each other. Situation awareness suffers when the operator's cognitive resources are exceeded; overloaded human information processing capacity isn't enough to process all available information.

Cognitive load

Cognitive load means the ratio of the operator's information processing capacity and mental demands imposed by a task. If cognitive load exceeds the capacity of the operator, the operator faces an overload situation. Overload exposes the operator to various performance errors and accidents over time (Phillips-Wren & Adya, 2020; Vidulich & Tsang, 2012). Managing cognitive load is especially important for decision-making and therefore task performance as well as avoiding errors. Cognitive overload limits performance whereas underload leads to impaired vigilance and inattentiveness, where important signals may go unnoticed (Hussein & Abbass, 2018).

Multitasking, which involves interruptions (e.g., alerts) and task switching, increases the operator's cognitive load and impairs situation awareness. Managing a swarm contains considerably more multitasking than for example controlling an individual drone (Hocraffer & Nam, 2017). When working with a swarm, the operator has to control or at least monitor several objects simultaneously. In addition to this, they have to be able to analyse large amounts of data provided by the swarm (e.g., what the swarm has done and observed) in different modalities and perspectives, which is taxing on attention and working memory. This can result in impaired situation awareness, mental fatigue, heightened risk for errors and slower decision-making (Hussein & Abbass, 2018; Källbäcker & Bjurling, 2023).

The number of active robots increases perceived task difficulty, especially when the interaction duration is long (Kaduk et al., 2024). More robots and longer task duration also result in higher emotional arousal, which could reflect increased cognitive demands. Thus, it should be taken into account that prolonged missions are likely to cause the operator stress and cognitive load.

The limitations of communication technology also increase cognitive load. If communication delays are substantial and data received possibly inaccurate, the operator needs to base their decisions on outdated or incomplete information (Hepworth et al., 2020). The need to estimate the difference between information displayed by the system and real values is cognitively demanding (P. Walker et al., 2024). Moreover, the operator and swarm do not act in a vacuum; usually the operator needs to communicate with other personnel, too, which imposes additional cognitive load (Dixon et al., 2005; P. Walker et al., 2024).

Estimating and predicting cognitive load

Being able to estimate the operator's cognitive state in real-time is very important with respect to developing adaptive systems. Adapting the level of autonomy according to the operator's cognitive load is especially important when over- or underload could cause errors with serious consequences (Hussein et al., 2022). Future swarm systems should be able to apraise the operator's performance, assess the risk for over- or underload, and adapt or delegate tasks accordingly to ensure the operator's ability to continue their activity optimally (Alharasees & Kale, 2024; Howard, 2007).

Cognitive load is traditionally measured with questionnaires, psychophysiological techniques (e.g. EEG, eye movements, pupil size, saccadic velocity, heart rate) and performance-based metrics (Zang et al., 2024). The most common self-report questionnaire, which usually forms the basis to estimating cognitive load, is NASA-TLX (Task Load Index) (Zang et al., 2024). The drawback is that the questionnaire is usually administered after the task, so it only offers very generalised estimates of cognitive load during the task. On the other hand, psychophysiological measurements are temporally more accurate but they often require expensive or complex equipment and a controlled test environment, which are not easy to implement in real-life operational tasks. Additionally, some psychophysiological methods are too sensitive to the subject's physical activity or environmental factors, which causes bias in the cognitive load estimates (Adams et al., 2023).

Machine learning and neural network models for estimating cognitive load have been widely studied in the context of driving (see e.g., He et al., 2022; Khan et al., 2024; Yoshida et al., 2014), but not so much regarding swarms. Zang et al. (2024) suggest a Large Language Model (LLM) based algorithm to replace traditional methods for estimating cognitive load in the context of swarms. This approach focuses on collecting information about the operator's actions and task performance using simple methods, which the LLM then analyses and generalises to various situations involving human-robot interaction. The specific model didn't yield particularly accurate results so far, but it still seems a possible method for predicting operator cognitive load and thus enhancing task performance in the future (Zang et al., 2024).

Managing cognitive load

Control methods. Control method choice affects the operator's cognitive load and swarm performance considerably. At its simplest, the operator can control a robot or a small group of "leaders" from which the effect of commands propagates to other members of the swarm (Kolling et al., 2015). In principle, this control method is similar to teleoperating a single robot, the only difference being that command effects spread through the whole swarm. A bit more refined idea of leader-based control has successfully been applied in simulated swarms consisting of as much as a hundred vehicles, as long as the operator focuses on groups sized not more than four robots at a time (Kerman et al., 2012). Leader-based control has been found to scale quite well to larger swarms, at least in information foraging tasks (Pendleton & Goodrich, 2013). A swarm can also be controlled indirectly by modifying its environment, either virtually or physically (Kolling et al., 2015). The operator can also modify the algorithms regulating the swarm's activity. However, to achieve desired outcomes, this control method requires extensive knowledge on the special characteristics of swarm behavior (Kolling et al., 2015).

Kolling et al. (2013) have compared leader-based and environmental control (n=32). The control method was either selecting subgroups of the swarm (leader) or a beacon with an attractor effect on robots at a specific range (environmental). Operators using the selection method, especially novices, performed better compared to those using beacons. The writers speculate, though, that beacon control might scale better to larger swarms as operating the beacon is less intense and involves fewer operator actions than the selection method (Kolling et al., 2013). With direct control the operator's role is more active, which can increase cognitive load (Hocraffer & Nam, 2017). This notion is supported by the finding that selecting individual robots or defining precise targets results in higher cognitive load when operating a large swarm (Adams et al., 2023). On the other hand, direct control enables more detailed and exact commands than indirect methods, which can be necessary for some missions. Nonetheless, in the light of current literature both selection and beacon methods have been found useful in operating swarms in complex environments (Kolling et al., 2015). It has also been proposed that combining leader-based and environmental control methods might reduce the operator's cognitive load compared to direct control only (P-tel et al., 2019).

Centralised, or leader-based control, is less efficient when the operator has limited or inaccurate information on the swarm's operational environment or in complex tasks which hinder the operator's ability to focus on all subtasks at the same time (Crandall et al., 2017). In these situations, centralised control in a way negates fault tolerance, the most essential strength of swarming. Crandall et al. (2017) suggest shared control as an efficient way to balance human control and the swarm's fault tolerance. In shared control, decision-making is divided between the operator and swarm dynamics. While preserving fault tolerance, the operator still has enough influence to manage the swarm according to mission objectives. This means that the operator conveys higher-level intentions through the control interface, whereas the robots take charge in lower-level decision-making in accordance with criteria defined beforehand (Crandall et al., 2017). In addition to preserving fault tolerance, the advantage of shared control is that it doesn't require a complete communication network between the operator and all members of the swarm – the operator commands parts of the swarm, and the information propagates through the whole swarm as a result of local interactions between swarm members (Soorati et al., 2021). As long as the operator is supplied with summarised information on the consensus variables of the swarm (location, movement direction), task performance remains at a high level (Nunnally et al., 2012). The usefulness of decentralised control is also highlighted in findings which report that operating in a peer-like manner from "inside" the swarm often yields more efficient results and helps avoid errors compared to a commander-like ("from above") control, especially in large swarms with high levels of autonomy (Hocraffer & Nam, 2017).

Shepherding. Traditional behavioral control methods, such as leader-follower control, are not necessarily optimal in, for example, military air surveillance missions; the leader vehicle in front faces a considerably higher risk of becoming neutralised when travelling toward a possibly dangerous environment. Air surveillance tasks also often require the ability to approach the target from several directions simultaneously, which isn't possible if all swarm members follow a single leader (Debie et al., 2021).

Asynchronous shepherding, inspired by sheep shepherding, can be used to face these challenges. For example in air surveillance missions, the operator acts as the shepherd commanding chosen robots from the swarm ("sheep dogs") to steer the rest of the swarm (Debie et al., 2021). An additional benefit of asynchronous shepherding is that not all robots in the swarm receive the command at the same time, but the commands rather propagate from one vehicle to another. This series of events can be visualised before actually launching the command, which gives the operator the opportunity to decline the chosen course of actions if it wouldn't result in a desired outcome (Debie et al., 2021).

Debie et al. (2021) have developed an automatic recommender system that recommends actions to the operator on how to carry out the shepherding task. The system was tested in a simulation environment. An artificial intelligence model (Human Factors Operating Picture (Rojas et al., 2020)) was created for monitoring the operator's performance and cognitive load in real-time. This AI component modifies the graphic interface based on the operator's cognitive needs, for example by estimating the optimal timing for displaying recommendations to the operator.

Shepherding has also been studied with digital twins, or creating digital copies of every swarm member (Nguyen et al., 2023). The operator commands virtual agents which in turn transfer the commands to the real swarm in the physical world. Virtual agents are cost-effective and enable easier implementation of high-level autonomy compared to a similar physical agent. This system was tested in a real-life case study in which the participants (n=2) were able to manage a swarm on two levels of autonomy. The performance was better on a higher level of autonomy.

Levels of autonomy. The level of autonomy (LOA) has a substantive effect on the operator's cognitive load (Hus-sein & Abbass, 2018). Rising the LOA lowers cognitive load in most cases. It is to be noted, though, that automation transfers workload to other tasks but doesn't remove it completely – a monitoring task can also be cognitively demanding (Mouloua & Hancock, 2019).

LOA can be categorised to four levels: full autonomy, machine-oriented semi-autonomy, human-oriented semi-autonomy and manual operation (Mi & Yang, 2013). A fully autonomous swarm completes its tasks without the operator's participation in decision-making. In machine-oriented semi-autonomy the swarm acts autonomously most of the time, relaying information to the operator on important events; in human-oriented semi-autonomy, the operator often instructs the swarm on decision-making. Manual operation means that the operator makes all decisions and commands all actions of the swarm.

On a low LOA, the operator is responsible for all low-level actions, which imposes a heavy cognitive load. It has been recommended that a higher LOA than manual operation should be preferred for swarms (Hocraffer & Nam, 2017). When the swarm operates fully or partially autonomously, the operator can treat the swarm as a single object which reduces cognitive load (Kolling et al., 2015). If lower-level actions, e.g. navigating, are executed autonomously by the swarm, the operator's cognitive capacity is freed for mission-level, monitoring work (Das et al., 2018). As an example, the Playbook™ approach employs plays or patterns defined to the swarm in advance, which can be called to play during a mission, much like how sports teams execute their actions (Miller et al., 2005). The plays consist of generalised rules according to which the swarm aims at accomplishing its tasks. When calling a play, the operator can issue more detailed task-related instructions or allow the swarm to autonomously choose the best course of actions within the limits of the play.

Raising the LOA enables managing larger swarms as monitoring is not as technologically heavy as teleoperating or controlling individual robots. With increased LOA, it is easier for the operator to monitor the whole swarm and maintain situation awareness (Amelink et al., 2008). However, (almost) full autonomy might impair the operator's situation awareness (Hussein & Abbass, 2018) if the information on what the swarm is doing and why is not easily accessible. It is also possible that when the swarm approaches full autonomy, the operator may begin to excessively trust the automation and thus experience a decay in their ability to react adequately in unexpected situations.

Compared to fully autonomous swarms and manual operation, it seems that the best results are achieved when the swarm and operator act in cooperation (Hocraffer & Nam, 2017). By comparing a swarm fully operated by a human to an autonomous swarm in the same task it can be investigated which tasks are better performed by the operator alone and which by the swarm automation. This information can be used to delegate tasks optimally both considering task performance and the operator's cognitive load (Hocraffer & Nam, 2017). It has been found that a human performs best in a monitoring role compared to a lower-level operating role (Hocraffer & Nam, 2017) which should be taken into account when developing operational environments (Hussein & Abbass, 2018).

Flexible autonomy. It is important to weigh the benefits gained from reducing the operator's cognitive load against the risks related to excessive trust and possible decline in the operator's situation awareness when planning the appropriate LOA for a swarm (Hocraffer & Nam, 2017). This challenge can be faced by employing flexible autonomy. Flexible autonomy means that the swarm's LOA is shifted according to the operator's needs and task demands (Chen & Barnes, 2014). The advantage of flexible autonomy is the ability to adapt to changes in real-time.

The initiative to change LOA can come from the human operator (adaptable) or the automation (adaptive). In adaptable systems, the operator decides on the appropriate LOA for each situation. The upside of this approach is that automation never replaces human as the decision-making authority. However, charging the operator with autonomy decisions might increase cognitive load and cause too much delay or complexity in decision-making, especially in already cognitively demanding situations (Chen & Barnes, 2014).

Adaptive systems, on the other hand, rely on the automation making LOA decisions based on criteria defined beforehand, for example according to the operator's cognitive load or environmental factor. Adaptive systems are useful in keeping the operator engaged in the task without the risk of cognitive overload – with risk of boredom the operator is tasked with more actions, whereas in high-load situations tasks are transferred to the automation, in which case the operator doesn't need to consume their information processing resources on deciding when to switch LOA. The LOA can be switched quickly because the decision doesn't require permission from the operator. The downside of adaptivity, however, is that the authority to make decisions about autonomy is surrendered to the automation, which might not be the best solution ethically or technologically (Chen & Barnes, 2014).

To combine the advantages and counter disadvantages of both approaches, mixed-initiative systems have been proposed. Mixed-initiative systems employ sharing the responsibility of decision-making between the operator and automation. These systems enable quick reacting to different situations while preserving the operator as part of decision-making (Hussein & Abbass, 2018). Flexible automation also supports designing control interfaces according to the operator's cognitive needs, for example via supporting the operator's information processing where needed (Chen & Barnes, 2014).

Autonomous agents. Dividing the workload between multiple operators decreases cognitive load as long as the roles are defined clearly (Hocraffer & Nam, 2017). This idea can be applied to make larger swarms more manageable to a single operator by utilising one or many intelligent agents (IA). IAs take responsibility in lower-level tasks, synthesising and communicating task-relevant information to the operator, as well as aiding the operator in decision-making and maintaining situation awareness (Wohleber et al., 2023). In this type of team the human and agent(s) need to form a relationship that capitalises both parties' strengths: the agent is tasked with aiding human decision-making whereas the human needs to be able to recognise possible errors committed by the agent and intervene when needed (Wohleber et al., 2023). In order for this to succeed, the operator must understand the limitations of the IA and the operational environment.

LaMonica et al. (2022) propose using OODA loop (Observe, Orient, Decide, Act) and MAD (Mission Actor Diagram) in enhancing human teaming with an IA. By modelling decision-making paths related to a task it is possible to identify steps where a human operator would benefit from the aid of an IA. Utilising this method can help reduce the operator's cognitive load as well as improve decision-making speed and accuracy (LaMonica et al., 2022). More specific instructions on how to use OODA loop and MAD are reported in the paper by La-Monica et al. (2022).

Trust

Both excessive trust and distrust in an autonomous agent is detrimental for task performance (Chen & Barnes, 2014). On the one hand, placing too much trust in automation can cause overreliance on it and, consequently, disastrous outcomes and accidents. On the other hand, distrust in automation can result in missing out on the benefits automation has to offer. A human might decline automation if the chosen LOA is too high compared to trust (Hussein & Abbass, 2018). Therefore, it is important to strive for calibrating the operator's trust in automation according to the actual reliability of the autonomous agent (Hussein & Abbass, 2018). The task performance of a robot has been found to be the main factor in determining trust (Hancock et al., 2011). Another important factor is transparency.

Transparency information includes information related to decision-making processes, predicted outcomes and uncertainty that is explicitly communicated by the autonomous agent to the operator. Traditionally it has been thought that transparency information provided by an agent may aid the operator's task performance and support trust calibration (Chen & Barnes, 2015; Stowers et al., 2017). The effects of transparency can however vary depending on the agent's actual abilities (Chen & Barnes, 2015) and communicating uncertainty information might sometimes impair the perceived usability of an agent as well as trust in it (Stowers et al., 2017). In accordance with this, the newest studies have highlighted that the relationship between transparency, trust and task performance is complex.

Transparency information can be described with three levels:

1. what is happening right now and what is the agent trying to do
2. why is the agent doing what it is doing
3. what can the operator expect to happen to the agent in the near future (3) (Chen et al., 2014).

A study by Stowers et al. (2017) (n=30+53) found that operators make the most correct decisions when an autonomous agent presents information on all three levels of transparency. Cognitive load and decision times remained the same even with additional information, or transparency levels. The operator's trust in the agent and usability experience also were at the highest when all levels of transparency were used. When uncertainty information was added on top of the transparency information, the operators were even more accurate in making their decisions while cognitive load remained at the same level than before. As a result of the increased information load, however, the operators' decision times were longer and the perceived usability of the autonomous agent decreased. Based on these results, Stowers et al. (2017) suggest that autonomous agents should also provide information on their decision-making processes and possible sources of uncertainty. Considering time-critical tasks, however, it might sometimes also be justifiable not to report uncertainty information, since it increases the operator's decision-making times (Stowers et al., 2017).

Mualla et al. (2020) (n=27) found similar results while researching the best way for autonomous agents to present transparency information to operators. They reported that adding explanations related to the agent's actions and decision-making processes enhances the operator's understanding on the agent's activity, but too detailed explanations impose too much cognitive load on the operator. Therefore, it's preferable that the agent's explanations be filtered (based on operator preferences or environmental conditions, for example) to ensure optimal operator performance (Mualla et al., 2020).

Wang et al. (2016) (n=160) studied how trust is affected by explanations provided by a robot. There were two types of explanations: confidence (what is the probability that the robot is correct) and observations (based on what observations is the decision made). Adding any type of explanations enhanced trust, decision-making and team performance especially when the robot's ability was low. Apparently, the explanations helped the participants calibrate their trust correctly. In other words, they learned to recognise situations when the robot should be trusted. This effect was not found when the robot's ability was high; it is probable that it is beneficial to trust a highly capable robot regardless of the explanations it might provide. Confidence- and observations-based explanations didn't differ from each other in the measurements studied. It is noteworthy, though, that confidence-based explanations will not help the operator diagnose or fix a malfunctioning robot, whereas observations-based explanations may provide tools for troubleshooting (Wang et al., 2016). P. Walker et al. (2024) also suggest that better diagnostic data would aid the operator in situations where automation has committed errors. Another study (n=65) yet highlights that communicating uncertainty information enhances trust in the robot (Kox et al., 2022).

Wohleber et al. (2023) (n=52) found a bit contradicting results on transparency information provided by an autonomous agent. The main finding was that the agent taking a criticising approach negatively affected the participants' willingness to heed to the recommendations of the agent, but only on a high level of transparency. Based on this, levels of transparency should be adjusted according to mission objectives and the agent itself. For example, a system whose main task is to scrutinise the operator's decisions and recommend improvements transparently, should avoid making the operator feel too criticised – this impairs trust calibration. On the other hand, if the developer knows their system is reliable but its decision-making processes are not easily explained, operator trust could be enhanced by adding some criticism to the agent's recommendations (Wohleber et al., 2023). In this study, verbal communication of transparency information was the best for low transparency, but graphic and iconographic approaches raised in popularity on a higher level of transparency, when the amount of information displayed was larger.

Wright et al. (2020) (n=56) studied the effect of an autonomous agent's transparency and reliability on the operator's performance, cognitive load, situation awareness and trust in the agent. Reliability and transparency did not affect cognitive load experienced by the operator in the studied task. It was found that the reliability of the agent had an effect on the operator's trust, but transparency did not. These results differ from those presented above. Errors made by the robot had a long-lasting effect on the operator's assumption of the robot's future trustworthiness. Errors led to lowered trust in the robot even when it didn't make any more errors. The trust recovered slowly over time if the agent remained reliable. These results should be taken into account when designing automated systems where the user will be aware of errors made by the system.

Interfaces

Because swarms are decentralised systems and technology limits real-time communication, usually only incomplete data on the state of the swarm is available. Operators often need to make decisions based on summarised data or imperfect observations (Mouloua & Hancock, 2019). For this reason it is extremely important to design interfaces that are intuitive and natural to the user, and that strive to display data as conveniently as possible. A well-designed interface can enhance system performance by taking advantage of things humans are good at while allowing the swarm to utilise its own strengths (Hocraffer & Nam, 2017). Comparably, an unfitting interface may demand too much from the operator's information processing capacity and/or result in errors and loss of situation awareness.

Cognitive fit. The main idea of cognitive fit is that the way information is presented has a direct effect on task performance. Representations better suited for a task result in better performance (Cabanag et al., 2024). A cognitively fit interface for one task isn't necessarily so for another task; an interface should always be designed according to the task in which it's going to be used. Generalised interfaces don't work as efficiently as interfaces specifically designed for a swarm or task (Lewis, 2013). Customisable interfaces also usually enable intuitiveness and lower cognitive load (Hocraffer & Nam, 2017).

Cabanag et al. (2024) give recommendations on designing cognitively fit interfaces:

- Perception. Effort related to perception should be minimised. The need for visual search should be minimised as well. Pre-attentive clues (clues that the brain processes very quickly, e.g., color, object size, opacity) should be utilised. Gestalt principles, such as the laws of similarity, proximity, continuity, closure, symmetry and organisation, should be leveraged.
- Counting. Mental operations should be minimised or externalised. It is useful to try to identify possibilities to let "visualisation do the counting" for the operator. This means that the way the data is visualised reduces the need to perform mental operations. For example, data to be compared should not be presented in a table but be embedded in pictures as some kind of perceptual elements, such as color. A perception task is much easier than a comparison task.
- Distraction. Distraction should be minimised. Salient cues and representations should be used. Especially acute or critical notifications should be salient so that the operator's attention is directed to them automatically (Szafir & Szafir, 2021). Presenting redundant or too much information should be avoided.
- Dynamic, operational contexts may include a lot of different tasks, and as stated above, different kinds of interfaces are a better fit for different tasks. It is not very reasonable to try to combine all appropriate elements from other tasks into one interface; then the amount of distraction will skyrocket. A better alternative could be some kind of an adaptive screen that would only display cognitively fit elements relevant to current active tasks.

The way that information is presented on an interface is a main factor in influencing operator's cognitive load. Too much information may exceed the operator's information processing capacity and complicate focusing on the essential, whereas too little information causes uncertainty and therefore cognitive load (Hussein et al., 2022). As a rule, information should be displayed at a rather high abstraction level while avoiding unnecessary details (Kolling et al., 2015; Szafir & Szafir, 2021). Instead of displaying the location and state of an individual robot, an overview of the states of the swarm, mission and environment should be preferred (Crandall et al., 2017; Hussein et al., 2022). Heat maps, for example, are useful when managing large swarms, in urgent tasks and in monitoring movement and area coverage (Soora-ti et al., 2021). The operator must however have the option to also examine more detailed information (individual swarm members) for e.g. troubleshooting (Amelink et al., 2008; Soorati et al., 2021; Szafir & Szafir, 2021). Predicting and visualising the future state of the swarm increases the operator's accuracy and decreases overcorrecting the swarm's behavior, especially in situations where the operator must act on incomplete data caused by high communication latency, for example (Kolling et al., 2015).

Seo et al. (2021) have studied how the information on robot states should be presented to make it as easy as possible to the operator. Most participants in the study reported that icons were more understandable than text. When the information was presented as text (1–2 words) the participants performed as well as with icons, but they felt the task was more difficult. This should be paid attention to if the operator's experience of cognitive load is essential for a task (Seo et al., 2021). Anthropomorphic representations, such as emojis, were seen as confusing. Color coding seemed like a good means to display the urgency level of robot states. Data is often visualised in red-green or rainbow colors (Szafir & Szafir, 2021). These aren't optimal choices, though – guidelines for efficient color coding are reported in a review by Silva et al. (2011). In addition to the aforementioned factors, size and contrast are important in visualising information (Xu et al., 2025).

When designing interfaces, the need for task switching should be minimised in order to maintain situation awareness at a high level (Hocraffer & Nam, 2017). Interruptions are especially detrimental during tasks that are cognitively demanding, such as planning and evaluation. Alerts, unavoidable as they are in operational environments, cause the operator to discontinue the ongoing task and force them to focus on details for a moment. Nevertheless, interfaces should be designed so that main tasks are only interrupted in critical situations or during low cognitive load (Chen & Barnes, 2014). It is important to encourage the maintenance of overview instead of attending to separate details (Bahodi et al., 2024) and to aid the operator in recovering situation awareness after an interruption (Lewis, 2013). Interfaces that enable higher-level monitoring seem to be most advantageous for performance (Lewis, 2013).

Interfaces customisable based on the operator's preferences result in lower cognitive load and should be favored (Hocraffer & Nam, 2017). Multimodal control mechanisms, such as touch screen and voice commands, may lower cognitive load by way of allowing the operator to choose a method that feels intuitive to them. Redundancy also adds an extra security measure in case the other system fails. Including several types of feedback (e.g., video, motion simulation, voice alerts) could also decrease cognitive load and improve situation awareness (Hocraffer & Nam, 2017). On the other hand, haptic feedback, for instance, might increase cognitive load in already challenging tasks but enhance performance when heads-up control is advantageous compared to a traditional interface (McDonald et al., 2017). First-person views should be avoided and instead use a third-person perspective, or a chase camera (Menda et al., 2011) if video feed is transmitted to the operator. It should be noted, though, that too much information is also detrimental; video feed sent from a large swarm probably increases the operator's cognitive load, even though when used appropriately, it could also improve situation awareness. It also seems that current technology at least does not enable receiving real-time video feed from a large swarm in a way that would be of significant benefit (Adams et al., 2023).

If decisions are made based on video feed sent by a swarm (e.g., identifying casualties) it is important to consider available resources (Abioye et al., 2023). High-quality images increase the operators' accuracy but also cognitive load as well as slow down decision-making due to the limits of human cognition and technology (especially latency). Analysing a low-quality image is faster for both the operators and communication technology, but results in accuracy declining a bit. It is noteworthy that humans are quite good at interpreting images with a lot of noise – it is probable, then, that low-quality images result in higher efficiency overall. There is a slightly increased risk of wrong calls, however (Abioye et al., 2023). If resources are very scarce, it might be more sensible to try and identify true positives on the first try, which would favor the slower but more accurate (high-quality) method (Abioye et al., 2023).

Controls. It has been thought that gesture-based interfaces might be intuitive and therefore reduce the operator's cognitive load. Gesture-based interfaces have been studied relatively much in the context of swarms. Traditionally these interfaces utilise a pre-programmed library of specific gestures that are mapped onto actions the swarm is capable of executing (Podevijn et al., 2013). This method enables for example selecting subgroups and steering them to different locations.

In the recent years, machine learning has drawn a lot of attention in studies focusing on gestures (Macchini et al., 2021; St-Onge et al., 2019; Suresh & Martínez, 2019). These studies have utilised wristbands that collect data on body positions and muscle activity. The data is relayed to a classifier system that learns the relationships between operator gestures and swarm actions by machine learning during a short calibration period. This enables the classifier to send commands to the swarm based on the gestures the operator performs. Models with machine learning allow for substantial flexibility because new gestures can be introduced quickly and all gestures can also easily be customised based on operator preferences. It has been shown that humans learn to use this kind of interfaces quickly and intuitively (Suresh & Martínez, 2019). The prolonged use of gestures might induce physical fatigue on the operator, though (Hou et al., 2017; Kim et al., 2024). Certain gestures and movement directions seem to be more demanding than others, so this should be taken into account when designing gesture-based user interfaces (Hou et al., 2017).

Other studied controls are for example gaming controllers (Da Silva Tchilian et al., 2020), tangible interfaces (Paas et al., 2022) and LLMs (Liu et al., 2025). There is some preliminary evidence of the functionality of these controls.

Applications of mixed reality (MR), augmented reality (AR) and virtual reality (VR) have been introduced recently in swarm research. Das et al. (2018) created an interface in MR by visualising information as easily interpretable 3D-holograms based on the priority level of the data. This is thought to decrease cognitive load since attention is directed to task-relevant data only. MR also enables easy implementation of the multimodal communication mentioned above (5.1.) (Das et al., 2018).

Xu et al. (2025) reason that the tablet-based interfaces that are currently in use increase cognitive load and impair situation awareness by forcing the operator to divide their attention between the tablet screen and real world. They studied (n=5) an AR heads-up interface that dynamically shifts between task- and safety views to avoid cognitive overload. Data relevant to the situation is displayed whereas irrelevant data is hidden. Compared to a regular tablet and a static AR interface, an adaptive AR interface lowered the operators' cognitive load and improved situation awareness without hindering task performance (Xu et al., 2025). It was found that the operators would benefit from clear instructions on when to assume manual control of the interface and when to switch back to autopilot. Visual cues, such as the future direction of the drones, would also aid in reorienting to the task view. This study also highlights the need for the interface to be easily customisable, as most of the studies presented above. In this study the operators only managed one drone at a time, but the writers propose that this interface could also be applied to swarms in situations where both reducing cognitive load and enhancing security are crucial.

Adams et al. (2023) used a VR based system (P. Walker et al., 2024) in their study. In this system the operator wears a VR device on their head and it provides a three-dimensional view of the situation. Two controllers were used in operating in the world. P. Walker et al. (2024) noticed that situation awareness was hindered by the narrowing of perspective in the VR environment caused by current VR technology; monitoring the environment is prevented while a detail is being inspected closely. Also unreliable automation and imprecise data from sensors compelled the operator to focus on smaller details instead of monitoring the big picture (P. Walker et al., 2024). The writers state that VR based interfaces should be designed realistically based on underlying technology since too utopistic plans will not succeed.

Case study: how large a swarm can one operator manage?

DARPA's (The Defense Advanced Research Projects Agency) OFFSET (OFFensive Swarm-Enabled Tactics) program tested how a single operator can monitor and manage a real-life heterogenous robot swarm in a complex mission in an authentic urban environment. This is the first study that has reported information on human performance as a swarm operator in a real-life operation (Adams et al., 2023).

In the study, one operator managed a swarm of a hundred robots. The operator led the swarm from a command base near the "battlefield", so communication latency was low. The operator did not have visuals on the swarm or the operational environment. The operator managed the swarm with plans created before and tactics created during the mission. The swarm was controlled with a VR-based Immersive Interaction Interface (3I) designed by CCAST (Command and Control of Aggregate Swarm Tactics). More detailed information on the system architecture of CCAST (Clark et al., 2021) and 3I (P. Walker et al., 2024) is reported in the mentioned papers.

Cognitive load experienced by the operator was measured by creating a multidimensional workload algorithm based on cognitive, speech, auditory and physical components that were combined with separate measurements of visual workload. Measurements were heart rate, heart rate variability, respiration rate, body posture, speech rate, voice intensity, voice activity, pitch and noise level. The algorithm was found to be sensitive to changes in workload levels, even though the model overestimated physical load because of the measurements used. The operator's workload seemed to be increased by creating tactics during the mission, especially when the operator couldn't rely on the automatic allocation provided by the system but had to select robots by themselves. Experience of overload was also generated when there was substantial delay in communications and telemetric data was outdated. Communicating with a (human) team member also felt burdensome because it requires interrupting the current task and switching from the virtual environment to the real world.

The estimates of operator cognitive load crossed the overload threshold a few times but remained within normal limits 96% of the time. The operator was able to succeed in tasks even in very challenging operational circumstances. It is to be noted, that there were only two participants in this study. They both were extremely fluent computer users and had been part of the system development team of this project. That is to say, they were exceptionally familiar with the equipment beforehand. The results qualitatively suggest that a single human's information processing capacity might be enough to control a heterogenous swarm of a hundred robots on a real-life mission (Adams et al., 2023). However, the results cannot be generalised to say that anyone could succeed in this kind of a mission without substantive training.

Recommendations on swarm operators' characteristics

Personal attributes play an important role in a swarm operator's ability to perform as efficiently as possible while maintaining situation awareness and avoiding cognitive overload. Appropriately designed interfaces might sometimes compensate individual differences (Rodes & Gugerty, 2012) but it's most likely that well-designed interfaces need to be combined with high ability on relevant skills in order to reach the best possible performance.

Swarm operator performance is most strongly predicted by multitasking ability (Brieber, 2025). Other particularly important factors are spatial skills, attentional control, video gaming experience (Chen & Barnes, 2012), working memory (especially visuospatial), directing attention (Hussein & Abbass, 2018), mathematical-logical reasoning, perceptual speed (Brieber, 2025), visual search and multi-object tracking (Memar & Esfahani, 2018). Stress tolerance, conscientiousness and emotional stability are personality characteristics that are related to better performance (Brieber, 2025).

The significance of attentional control is highlighted in challenging multitasking environments and when an autonomous agent is miss-prone, which means that they have the tendency to ignore situations where actions are needed (Chen & Barnes, 2012). Visual search and multi-object tracking predict better cooperation with an AI agent whose job is to provide feedback and suggest courses of action – operators with lower ability are more likely to ignore information displayed by an agent, whereas high-ability individuals are able to process the information and make advised decisions based on it (Distefano, 2025). Video gaming experience as well has a positive impact on cooperation with an autonomous agent (Chen & Barnes, 2012).

At least some skills mentioned above can be improved with training. Video games can be played at any time, of course, and benefits can be found even for fairly short training. Multi-tasking and related attentional regulation skills can be trained with gaze-training tasks (Wilson et al., 2011). There is also preliminary evidence (n=12) that AI assistance could aid in lowering cognitive load while a person practices motor multitasking skills compared to training without AI assistance (Talypova et al., 2024). The same study also found that the skills learned during training were preserved even though the AI assistance was later removed. This kind of a training method might be applicable to the multitasking-heavy swarm operating context.

Task experience matters, too: novices experience more cognitive load in managing a swarm than experts, especially when the swarm's LOA is low (Alharasees & Kale, 2024; Hussein & Abbass, 2018). The most prominent challenge for novices in understanding and predicting swarm actions is that swarm behavior is often unintuitive to humans (Hocraffer & Nam, 2017). Two things are particularly important in successfully managing a robot swarm: understanding of the swarm's state and the ability to predict what kind of behavior the input given by an operator will produce (Nagavalli et al., 2015). This means that the operator needs to develop an accurate mental model of swarm behavior through experience.

A phenomenon called neglect benevolence has been found to affect swarms. Neglect benevolence means that for the swarm to function optimally, it's beneficial that the operator doesn't send commands to the swarm at too short intervals, but rather "neglects" the swarm for a while after each input. If commands are given too rapidly, the swarm algorithms don't have time to stabilise between the inputs, which might result in surprising and undesired outcomes (P. Walker et al., 2012). However, too long intervals between inputs might lead to poor performance as well, so it's important that the neglect interval stay within certain limits. Through learning, humans develop an ability to estimate the optimal timing for sending commands to a swarm (Nagavalli et al., 2015). Interfaces can be embedded with visual aids that help the operator understand how the swarm configuration will look like after a command. Naïve participants perform better with an aided interface. This points to the theory that the challenges in human-swarm interaction emerge from the counterintuitiveness of swarm behavior and movement, and performance improves when the variables are displayed in a perceptually accessible way (Nagavalli et al., 2015). Automated aids that use visualisation to help the operator understand the swarm's state could be beneficial during swarm operator training (Capiola et al., 2022; Nagavalli et al., 2015).

It is also worth noting that VR environments inflict nausea in many people (Chattha et al., 2020). If interfaces are heavily based on virtual reality, this should be taken into consideration when choosing swarm operators.

Future views

Ukraine's Operation Spider's Web showed that careful mission planning can help achieve tremendous results with only comparably small expenses. Planning the operation took 18 months, though (Gozzi & Verify, 2025). Such preparation times are not very efficient in fast-paced warfare. So, it would be important to develop swarm systems to be easily utilised on a wide variety of missions at a quick schedule. It cannot be too much emphasised that customisable interfaces and swarm control methods play a key role in this.

Most of swarm research has been conducted in simulation environments. It is thus still somewhat difficult to predict how humans will actually interact with swarms in real-life. Experiences in real-life might reveal some surprising challenges that aren't realised in simulations. It has for example already been noticed, that humans' psychophysiological reactions in a real environment grow significantly stronger with increasing number of robots compared to a simulation (Kaduk et al., 2024; Podevijn, O'Grady, Mathews, et al., 2016). This is why it should be a priority to invest in research conducted with real robots. If for some reason it is not possible to use real robots, virtual reality should prove a better choice than a simulation on a computer screen (Podevijn, O'Grady, Fantini-Hauwel, et al., 2016).

Autonomy in simulations and real-life might also not be fully corresponding to each other. Higher levels of autonomy are considerably easier to simulate than to implement in real robots, at least for now. At the same time, it must be remembered that research consistently takes great leaps at advancing artificial intelligence. Because AI develops so fast, system designers shouldn't persist on some model that seems the best right now. Our systems need to be flexible enough to be able to incorporate and utilise new innovations that are bound to emerge.

A major factor in slowing down the actualisation of adaptive systems is the inability to reliably measure or model the operator's cognitive load during an authentic mission (Rojas et al., 2020).

Currently, measuring cognitive load requires either impractical and expensive equipment or is based on inaccurate methods. It would therefore be important to invest in research on modelling the operator's cognitive load. If a sufficiently accurate but easily implementable method to estimate cognitive load can be created, the full potential of adaptive systems or systems with flexible autonomy can be unlocked.

Ethical questions are of utmost importance considering weapons systems that are even partially autonomous. Who makes operational decisions? Do situations exist where full autonomy can be given to a machine? How can we be sure that an autonomous system is reliable and predictable and that it will make decisions that are understandable to humans?

Conclusion

The major human factors challenges in human-swarm interaction are situation awareness, cognitive load and trust in automation. Situation awareness has a considerable impact on decision-making and task performance. Situation awareness is hindered by communication challenges, inappropriate interfaces and cognitive overload. Cognitive load should not be too low or too high, since both cause errors and impair decision-making. Developing a method to measure cognitive load is an important goal considering systems that utilise flexible autonomy. Cognitive load is caused by multitasking, larger swarms, manual operations performed by the operator, information overflow and communication challenges. Cognitive load can be managed by choosing control methods and interfaces best suited for a specific swarm, by increasing the swarm's automatism and by using autonomous AI agents to support decision-making. Trust in an autonomous agent must be correctly calibrated. Both trusting too much and too little is suboptimal. Factors influencing trust are e.g. the actual reliability of an agent and transparency information provided to the operator.

Interfaces must be designed according to the principles of cognitive fit. Using technology should be as intuitive and effortless as possible, so that the operator can focus on the main tasks. Interfaces should be designed to be easily customisable based on the task, swarm or operator preferences. Gesture- and VR-based interfaces seem promising, but they also involve new challenges, such as muscle fatigue or VR-induced motion sickness.

Initial empirical results show that it is possible for one person to manage a heterogenous swarm as large as a hundred robots in optimal circumstances. Important abilities and personality characteristics for swarm operators are multitasking, spatial skills, attention, working memory, mathematical-logical reasoning, perceptual speed, stress tolerance, conscientiousness and emotional stability.

In the future, research efforts should be focused on developing suitable interfaces, modelling cognitive load and conducting experiments in real-life settings instead of simulations.

For More Information

PhD (Psychology), adjunct professor Lauri Oksama is a Chief Scientist at Finnish Defence Research Agency.

References

Abioye, A. O., Naiseh, M., Hunt, W., Clark, J., Ramchurn, S. D., & Soorati, M. D. (2023). The Effect of Data Visualisation Quality and Task Density on Human-Swarm Interaction. 2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 1494–1501. https://doi.org/10.1109/RO-MAN57019.2023.10309454

Adams, J. A., Hamell, J., & Walker, P. (2023). Can a single human supervise a swarm of 100 heterogeneous robots? Field Robotics, 3, 837–881.

Alharasees, O., & Kale, U. (2024). Human Factors and AI in UAV Systems: Enhancing Operational Efficiency Through AHP and Real-Time Physiological Monitoring. Journal of Intelligent & Robotic Systems, 111(1), 5. https://doi.org/10.1007/s10846-024-02188-y

Amelink, M. H., Mulder, M., & Van Paassen, M. (2008). Designing for human-automation interaction: Abstraction-sophistication analysis for UAV control. Proceedings of the International Multiconference of Engineers and Computer Scientists (IMECS), 2168(1), 19–21.

Bahodi, M.-T., van Berkel, N., Skov, M., & Merritt, T. (2024). Show Me What's Wrong: Impact of Explicit Alerts on Novice Supervisors of a Multi-Robot Monitoring System. Proceedings of the Second International Symposium on Trustworthy Autonomous Systems. https://doi.org/10.1145/3686038.3686069

Bjurling, O. (2025). Designing Human-Swarm Interaction Systems (Vol. 900). Linköping University Electronic Press.

Bondar, K. (2025). How Ukraine's Operation 'Spider's Web'Redefines Asymmetric Warfare. Center for Strategic and International Studies. February, 6.

Brieber, D. (2025). Psychological Aptitude for Drone Pilots: Key Traits and Cognitive Demands. Military psychology conference 2025.

Cabanag, M., Stanton, C. J., & Cass, J. R. (2024). Fighting Information Overload: Modelling, Predicting and Inducing Cognitive Fit in Human-Autonomy Teaming. Proceedings of the 12th International Conference on Human-Agent Interaction, 204–213. https://doi.org/10.1145/3687272.3688305

Capiola, A., Hamdan, I. aldin, Fox, E. L., Lyons, J. B., Sycara, K., & Lewis, M. (2022). "Is something amiss?" Investigating individuals' competence in estimating swarm degradation. Theoretical Issues in Ergonomics Science, 23(5), 562–587.

Chattha, U. A., Janjua, U. I., Anwar, F., Madni, T. M., Cheema, M. F., & Janjua, S. I. (2020). Motion sickness in virtual reality: An empirical evaluation. IEEE Access, 8, 130486–130499.

Chen, J. Y., & Barnes, M. J. (2012). Supervisory control of multiple robots: Effects of imperfect automation and individual differences. Human Factors, 54(2), 157–174.

Chen, J. Y., & Barnes, M. J. (2014). Human–agent teaming for multirobot control: A review of human factors issues. IEEE Transactions on Human-Machine Systems, 44(1), 13–29.

Chen, J. Y., & Barnes, M. J. (2015). Agent Transparency for Human-Agent Teaming Effectiveness. 2015 IEEE International Conference on Systems, Man, and Cybernetics, 1381–1385. https://doi.org/10.1109/SMC.2015.245

Chen, J. Y., Procci, K., Boyce, M., Wright, J. L., Garcia, A., & Barnes, M. (2014). Situation awareness-based agent transparency.

Clark, S., Usbeck, K., Diller, D., & Schantz, R. E. (2021). CCAST: A framework and practical deployment of heterogeneous unmanned system swarms. GetMobile: Mobile Computing and Communications, 24(4), 17–26.

Crandall, J. W., Anderson, N., Ashcraft, C., Grosh, J., Henderson, J., McClellan, J., Neupane, A., & Goodrich, M. A. (2017). Human-swarm interaction as shared control: Achieving flexible fault-tolerant systems. International Conference on Engineering Psychology and Cognitive Ergonomics, 266–284.

Cummings, M. L., & Mitchell, P. J. (2008). Predicting controller capacity in supervisory control of multiple UAVs. IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans, 38(2), 451–460.

Da Silva Tchilian, R., Moreno, U. F., & Netto, M. (2020). Assisted Teleoperation for a Human-Swarm Interaction System. IFAC PapersOnLine, 53(5), 602–607.

Das, A., Kol, P., Lundberg, C., Doelling, K., Sevil, H. E., & Lewis, F. (2018). A Rapid Situational Awareness Development Framework for Heterogeneous Manned-Unmanned Teams. NAECON 2018 - IEEE National Aerospace and Electronics Conference, 417–424. https://doi.org/10.1109/NAECON.2018.8556769

Debie, E., El-Fiqi, H., Fidock, J., Barlow, M., Kasmarik, K., Anavatti, S., Garratt, M., & Abbass, H. (2021). Autonomous recommender system for reconnaissance tasks using a swarm of UAVs and asynchronous shepherding. Human-Intelligent Systems Integration, 3(2), 175–186. https://doi.org/10.1007/s42454-020-00024-w

Distefano, J. P. (2025). Augmenting artificial intelligence design and performance with human physiological and behavioral information. Dissertation Abstracts International: Section B: The Sciences and Engineering, 86(3-B), No-Specified. APA PsycInfo <2024>.

Dixon, S. R., Wickens, C. D., & Chang, D. (2005). Mission control of multiple unmanned aerial vehicles: A workload analysis. Human Factors, 47(3), 479–487.

Elkin-Frankston, S., Horner, C., Alzahabi, R., & Cain, M. S. (2023). Characterizing motion prediction in small autonomous swarms. Applied Ergonomics, 106, 103909. https://doi.org/10.1016/j.apergo.2022.103909

Endsley, M. R. (1988). Situation awareness global assessment technique (SAGAT). Proceedings of the IEEE 1988 National Aerospace and Electronics Conference, 789–795.

Endsley, M. R., & Garland, D. J. (2000). Situation awareness analysis and measurement. CRC press.

Gozzi, L., & Verify, B. (2025). How Ukraine carried out daring ‘Spider Web’attack on Russian bombers. BBC News.

Hancock, P. A., Billings, D. R., Schaefer, K. E., Chen, J. Y., De Visser, E. J., & Parasuraman, R. (2011). A meta-analysis of factors affecting trust in human-robot interaction. Human Factors, 53(5), 517–527.

He, D., Wang, Z., Khalil, E. B., Donmez, B., Qiao, G., & Kumar, S. (2022). Classification of driver cognitive load: Exploring the benefits of fusing eye-tracking and physiological measures. Transportation Research Record, 2676(10), 670–681.

Hepworth, A. J., Baxter, D. P., Hussein, A., Yaxley, K. J., Debie, E., & Abbass, H. A. (2020). Human-swarm-teaming transparency and trust architecture. IEEE/CAA Journal of Automatica Sinica, 8(7), 1281–1295.

Hocraffer, A., & Nam, C. S. (2017). A meta-analysis of human-system interfaces in unmanned aerial vehicle (UAV) swarm management. Applied Ergonomics, 58, 66–80. https://doi.org/10.1016/j.apergo.2016.05.011

Hou, W., Wu, C., & Chen, X. (2017). The research of wearable device user fatigue based on gesture interaction. International Conference on Human-Computer Interaction, 172–182.

Howard, A. M. (2007). A Systematic Approach to Predict Performance of Human–Automation Systems. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 37(4), 594–601. https://doi.org/10.1109/TSMCC.2007.897505

Hussein, A., & Abbass, H. (2018). Mixed Initiative Systems for Human-Swarm Interaction: Opportunities and Challenges. 2018 2nd Annual Systems Modelling Conference (SMC), 1–8. https://doi.org/10.1109/SYSMC.2018.8509744

Hussein, A., Ghignone, L., Nguyen, T., Salimi, N., Nguyen, H., Wang, M., & Abbass, H. A. (2022). Characterization of Indicators for Adaptive Human-Swarm Teaming. https://research.ebsco.com/linkprocessor/plink?id=e833672f-d603-3ac8-99b4-3836cde08c30

Kaduk, J., Cavdan, M., Drewing, K., & Hamann, H. (2024). From One to Many: How Active Robot Swarm Sizes Influence Human Cognitive Processes. 2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN), 1207–1212. https://doi.org/10.1109/RO-MAN60168.2024.10731232

Kaduk, J., Cavdan, M., Drewing, K., Vatakis, A., & Hamann, H. (2023). Effects of Human-Swarm Interaction on Subjective Time Perception: Swarm Size and Speed. Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction, 456–465. https://doi.org/10.1145/3568162.3578626

Källbäcker, J., & Bjurling, O. (2023). Human-Swarm Interaction in Semi-voluntary Search and Rescue Operations: Opportunities and Challenges. Proceedings of the European Conference on Cognitive Ergonomics 2023. https://doi.org/10.1145/3605655.3605684

Kerman, S., Brown, D., & Goodrich, M. A. (2012). Supporting human interaction with robust robot swarms. 2012 5th International Symposium on Resilient Control Systems, 197–202.

Khan, M. A., Asadi, H., Qazani, M. R. C., Arogbonlo, A., Nahavandi, S., & Lim, C. P. (2024). Predicting cognitive load in immersive driving scenarios with a hybrid CNN-RNN model. International Conference on Neural Information Processing, 224–240.

Kim, J. H., Pulipati, V., Yu, C.-Y., & Wang, F. (2024). Understanding the influence of fatigue on full arm gestures in augmented reality environments. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 68(1), 1194–1199.

Kolling, A., Sycara, K., Nunnally, S., & Lewis, M. (2013). Human swarm interaction: An experimental study of two types of interaction with foraging swarms. Journal of Human-Robot Interaction, 2(2).

Kolling, A., Walker, P., Chakraborty, N., Sycara, K., & Lewis, M. (2015). Human interaction with robot swarms: A survey. IEEE Transactions on Human-Machine Systems, 46(1), 9–26.

Kox, E. S., Siegling, L. B., & Kerstholt, J. H. (2022). Trust Development in Military and Civilian Human–Agent Teams: The Effect of Social-Cognitive Recovery Strategies. International Journal of Social Robotics, 14(5), 1323–1338. https://doi.org/10.1007/s12369-022-00871-4

LaMonica, D. A., Drnec, K., & Miller, M. E. (2022). Employing MBSE to Assess and Evaluate Human Teaming in Military Aviation Command and Control. 2022 IEEE 3rd International Conference on Human-Machine Systems (ICHMS), 1–7. https://doi.org/10.1109/ICHMS56717.2022.9980801

Lewis, M. (2013). Human interaction with multiple remote robots. Reviews of Human Factors and Ergonomics, 9(1), 131–174.

Liu, J., Li, H., Chai, C., Chen, K., & Wang, D. (2025). A LLM-informed multi-agent AI system for drone-based visual inspection for infrastructure. Advanced Engineering Informatics, 68, 103643. https://doi.org/10.1016/j.aei.2025.103643

Macchini, M., De Matteïs, L., Schiano, F., & Floreano, D. (2021). Personalized human-swarm interaction through hand motion. IEEE Robotics and Automation Letters, 6(4), 8341–8348.

McDonald, S. J., Colton, M. B., Alder, C. K., & Goodrich, M. A. (2017). Haptic Shape-Based Management of Robot Teams in Cordon and Patrol. https://research.ebsco.com/linkprocessor/plink?id=5d681669-3113-37a4-85f6-3092bacacabe

Memar, A. H., & Esfahani, E. T. (2018). Physiological measures for human performance analysis in human-robot teamwork: Case of tele-exploration. IEEE Access, 6, 3694–3705.

Menda, J., Hing, J. T., Ayaz, H., Shewokis, P. A., Izzetoglu, K., Onaral, B., & Oh, P. (2011). Optical brain imaging to enhance UAV operator training, evaluation, and interface development. Journal of Intelligent & Robotic Systems, 61(1), 423–443.

Mi, Z.-Q., & Yang, Y. (2013). Human-robot interaction in UVs swarming: A survey. International Journal of Computer Science Issues (IJCSI), 10(2 Part 1), 273.

Miller, C., Funk, H., Wu, P., Goldman, R., Meisner, J., & Chapman, M. (2005). The PlaybookTM approach to adaptive automation. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 49(1), 15–19.

Mouloua, M., & Hancock, P. A. (2019). Human Performance in Automated and Autonomous Systems: Emerging Issues and Practical Perspectives. CRC Press.

Mualla, Y., Tchappi, I. H., Najjar, A., Kampik, T., Galland, S., & Nicolle, C. (2020). Human-agent Explainability: An Experimental Case Study on the Filtering of Explanations. ICAART (1), 378–385.

Nagavalli, S., Chien, S.-Y., Lewis, M., Chakraborty, N., & Sycara, K. (2015). Bounds of Neglect Benevolence in Input Timing for Human Interaction with Robotic Swarms. Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, 197–204. https://doi.org/10.1145/2696454.2696470

Nguyen, H., Hussein, A., Garratt, M. A., & Abbass, H. A. (2023). Swarm Metaverse for Multi-Level Autonomy Using Digital Twins. https://research.ebsco.com/linkprocessor/plink?id=115ffed9-1235-3ef5-87cd-4d71f44f62f6

Nisser, T., & Westin, C. (2006). Human factors challenges in unmanned aerial vehicles (uavs): A literature review. School of Aviation of the Lund University, Ljungbyhed, 50.

Nunnally, S., Walker, P., Kolling, A., Chakraborty, N., Lewis, M., Sycara, K., & Goodrich, M. (2012). Human influence of robotic swarms with bandwidth and localization issues. https://research.ebsco.com/linkprocessor/plink?id=dbe528ac-45f6-370b-9a1a-354e19d7584a

Özyörük, H. E. (2020). Systematic analysis and classification of the literature regarding the impact of human factors on unmanned aerial vehicles (UAV). Journal of Aviation, 4(2), 71–81.

Paas, A., Coffey, E. B. J., Beltrame, G., & St-Onge, D. (2022). Towards evaluating the impact of swarm robotic control strategy on operators' cognitive load. 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 217–223. https://doi.org/10.1109/RO-MAN53752.2022.9900763

Parush, A., Kramer, C., Foster-Hunt, T., Momtahan, K., Hunter, A., & Sohmer, B. (2011). Communication and team situation awareness in the OR: Implications for augmentative information display. Journal of Biomedical Informatics, 44(3), 477–485.

Patel, J., Xu, Y., & Pinciroli, C. (2019). Mixed-Granularity Human-Swarm Interaction. 2019 International Conference on Robotics and Automation (ICRA), 1059–1065. https://doi.org/10.1109/ICRA.2019.8793261

Pendleton, B., & Goodrich, M. (2013). Scalable human interaction with robotic swarms. AIAA Infotech@ Aerospace (I@a) Conference, 4731.

Phillips-Wren, G., & Adya, M. (2020). Decision making under stress: The role of information overload, time pressure, complexity, and uncertainty. Journal of Decision Systems, 29(sup1), 213–225.

Podevijn, G., O'Grady, R., Fantini-Hauwel, C., & Dorigo, M. (2016). Investigating the effect of the reality gap on the human psychophysiological state in the context of human-swarm interaction. https://research.ebsco.com/linkprocessor/plink?id=3389c6bd-1789-3835-80ac-6ce32a66d7da

Podevijn, G., O'Grady, R., Mathews, N., Gilles, A., Fantini-Hauwel, C., & Dorigo, M. (2016). Investigating the effect of increasing robot group sizes on the human psychophysiological state in the context of human–swarm interaction. Swarm Intelligence, 10(3), 193–210.

Podevijn, G., O'Grady, R., Nashed, Y. S., & Dorigo, M. (2013). Gesturing at subswarms: Towards direct human control of robot swarms. Conference Towards Autonomous Robotic Systems, 390–403.

Rodes, W., & Gugerty, L. (2012). Effects of electronic map displays and individual differences in ability on navigation performance. Human Factors, 54(4), 589–599.

Rojas, R. F., Debie, E., Fidock, J., Barlow, M., Kasmarik, K., Anavatti, S., Garratt, M., & Abbass, H. A. (2020). Human performance operating picture for shepherding a swarm of autonomous vehicles. Shepherding UxVs for Human-Swarm Teaming: An Artificial Intelligence Approach to Unmanned X Vehicles, 293–323.

Seo, S. H., Young, J. E., & Irani, P. (2021). How are Your Robot Friends Doing? A Design Exploration of Graphical Techniques Supporting Awareness of Robot Team Members in Teleoperation. International Journal of Social Robotics, 13(4), 725–749. https://doi.org/10.1007/s12369-020-00670-9

Shappell, S. A., & Wiegmann, D. A. (2000). The human factors analysis and classification system–HFACS.

Silva, S., Santos, B. S., & Madeira, J. (2011). Using color in visualization: A survey. Computers & Graphics, 35(2), 320–333.

Soorati, M. D., Clark, J., Ghofrani, J., Tarapore, D., & Ramchurn, S. D. (2021). Designing a User-Centered Interaction Interface for Human–Swarm Teaming.

https://research.ebsco.com/linkprocessor/plink?id=ddd7edd5-5efc-3118-a537-fb371c85c932

St-Onge, D., Côté-Allard, U., Glette, K., Gosselin, B., & Beltrame, G. (2019). Engaging with Robotic Swarms: Commands from Expressive Motion. J. Hum.-Robot Interact., 8(2). https://doi.org/10.1145/3323213

Stowers, K., Kasdaglis, N., Rupp, M., Chen, J., Barber, D., & Barnes, M. (2017). Insights into Human-Agent Teaming: Intelligent Agent Transparency and Uncertainty. In P. Savage-Knepshield & J. Chen (Eds.), Advances in Human Factors in Robots and Unmanned Systems (Vol. 499, pp. 149–160). Springer International Publishing. https://doi.org/10.1007/978-3-319-41959-6_13

Suresh, A., & Martínez, S. (2019). Gesture based Human-Swarm Interactions for Formation Control using interpreters. IFAC PapersOnLine, 51(34), 83–88.

Szafir, D., & Szafir, D. A. (2021). Connecting human-robot interaction and data visualization. Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, 281–292.

Talypova, D., Lingler, A., & Wintersberger, P. (2024). Prolonged Usage of AI Assistant for Improving Multitasking Performance. Proceedings of the 12th International Conference on Human-Agent Interaction, 404–407. https://doi.org/10.1145/3687272.3690898

Vidulich, M. A., & Tsang, P. S. (2012). Mental workload and situation awareness. Handbook of Human Factors and Ergonomics, 243–273.

Walker, C. (2025, June 4). Drone attack on Russian bombers costs Putin £5bn—But it's a lot worse than that. Express. https://www.express.co.uk/news/world/2064188/operation-spiders-web-russia-damage

Walker, P., Hamell, J., Miller, C., Ladwig, J., Wauck, H., & Keller, P. (2024). Immersive interaction interface (I3): A virtual reality swarm control interface. IEEE Transactions on Field Robotics.

Walker, P., Nunnally, S., Lewis, M., Kolling, A., Chakraborty, N., & Sycara, K. (2012). Neglect benevolence in human control of swarms in the presence of latency. 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC), 3009–3014.

Wang, N., Pynadath, D. V., & Hill, S. G. (2016). Trust Calibration within a Human-Robot Team: Comparing Automatically Generated Explanations. The Eleventh ACM/IEEE International Conference on Human Robot Interaction, 109–116.

Wilson, M. R., Vine, S. J., Bright, E., Masters, R. S., Defriend, D., & McGrath, J. S. (2011). Gaze training enhances laparoscopic technical skill acquisition and multi-tasking performance: A randomized, controlled study. Surgical Endoscopy, 25(12), 3731–3739.

Wohleber, R. W., Stowers, K., Barnes, M., & Chen, J. Y. C. (2023). Agent transparency in mixed-initiative multi-UxV control: How should intelligent agent collaborators speak their minds? Computers in Human Behavior, 148, 107866. https://doi.org/10.1016/j.chb.2023.107866

Wright, J. L., Chen, J. Y. C., & Lakhmani, S. G. (2020). Agent Transparency and Reliability in Human–Robot Interaction: The Influence on User Confidence and Perceived Reliability. IEEE Transactions on Human-Machine Systems, 50(3), 254–263. https://doi.org/10.1109/THMS.2019.2925717

Xu, P., Garcia, J., Ooi, W. T., & Jouffrais, C. (2025). SafeSpect: Safety-First Augmented Reality Heads-up Display for Drone Inspections. Proceedings of the 2025 CHI Conference on Hu-man Factors in Computing Systems. https://doi.org/10.1145/3706598.3714283

Yoshida, Y., Ohwada, H., Mizoguchi, F., & Iwasaki, H. (2014). Classifying cognitive load and driving situation with machine learning. International Journal of Machine Learning and Computing, 4(3), 210.

Zang, W., Hu, M., & Liu, R. (2024). Large Language Model Driven Interactive Learning for Real-Time Cognitive Load Prediction in Human-Swarm Systems. 2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN), 97–102. https://doi.org/10.1109/RO-MAN60168.2024.10731286