



HAL
open science

Towards Augmented Reality Support for Swarm Monitoring: Evaluating Visual Cues to Prevent Fragmentation

Aymeric Hénard, Étienne Peillard, Jérémy Rivière, Sébastien Kubicki, Gilles
Coppin

► **To cite this version:**

Aymeric Hénard, Étienne Peillard, Jérémy Rivière, Sébastien Kubicki, Gilles Coppin. Towards Augmented Reality Support for Swarm Monitoring: Evaluating Visual Cues to Prevent Fragmentation. IEEE Transactions on Visualization and Computer Graphics, In press, <10.1109/TVCG.2025.3616840>. <hal-05245362>

HAL Id: hal-05245362

<https://hal.science/hal-05245362v1>

Submitted on 8 Sep 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Towards Augmented Reality Support for Swarm Monitoring: Evaluating Visual Cues to Prevent Fragmentation

Aymeric Hénard , Étienne Peillard , Jérémy Rivière , Sébastien Kubicki , Gilles Coppin 




Fig. 1: (a) The baseline visualisation, where only the robots are visible; and the three proposed visualisations, (b) robots and links between interacting robots using grey lines of varying thickness, (c) robots and their dominant force using coloured arrows, (d) robots and their current direction using pink directional arrows

Abstract— Swarm fragmentation, the breakdown of communication and coordination among robots, can critically compromise a swarm’s mission. Integrating Augmented Reality support into swarm monitoring—especially through co-located visualisations anchored directly on the robots— may enable human operators to detect early signs of fragmentation and intervene effectively. In this work, we propose three localised visual cues—targeting robot connectivity, dominant decision influences, and movement direction—to make explicit the underlying Perception-Decision-Action (PDA) loop of each robot. Through an immersive Virtual Reality user study, 51 participants were tasked with both anticipating potential fragmentation and selecting the appropriate control to prevent it, while observing swarms exhibiting expansion, densification, flocking, and swarming behaviours. Our results reveal that a visualisation emphasising inter-robot connectivity significantly improves anticipation of fragmentation, though none of the cues consistently enhance control selection over a baseline condition. These findings underscore the potential of co-located AR-enhanced visual feedback to support human-swarm interaction and inform the design of future AR-based supervisory systems for robot swarms. A free copy of this paper and all supplemental materials are available at <https://osf.io/49gny>.

Index Terms— Augmented Reality, Virtual Reality, Human-Swarm Interaction, Robot Swarms, Visualisation, Supervision, User Study, Perception-Decision-Action Loop

1 INTRODUCTION

Swarms of robots are autonomous systems composed of low-capability robots [42] that exhibit self-organised collective behaviours, such as flocking [37], which require the coordination of the robots that compose them. However, robots’ technical limitations only allow them to interact over limited distances [9]. Some robots may cease to interact, and the swarm may split up. This phenomenon is known as swarm fragmentation [31]. As a result, the swarm may no longer be able to achieve its objective because part of itself is missing (e.g. coverage of a large area requiring a minimum number of robots). A solution to avoid fragmentation is to allow human users to monitor the swarm and intervene if necessary, as they can determine if fragmentation is detrimental to the swarm. Depending on the level of intervention, humans can take on different roles in interacting with a swarm [43]. As a supervisor, a human provides high-level commands and long-term objectives, such as assigning a rendezvous point, changing the swarm’s behaviour, or

• Aymeric Hénard, Étienne Peillard and Gilles Coppin are with IMT Atlantique, CNRS, Lab-STICC, Brest, France. E-mail: firstname.lastname@imt-atlantique.fr

• Jérémy Rivière is with Univ Brest, CNRS, Lab-STICC, Brest, France. E-mail: jriviere@univ-brest.fr

• Sébastien Kubicki is with ENIB, CNRS, Lab-STICC, Brest, France. E-mail: kubicki@enib.fr

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

selecting a subset of robots to patrol a specific area. In contrast, an operator focuses on low-level, short-term actions, like placing virtual beacons or adjusting parameters of the behaviour model, to influence the swarm. This role is more involved in monitoring and compensating for limitations in the swarm's autonomy, such as fragmentation. In this article, we focus on the operator role. Previous studies [5, 6, 15] have shown that operators with no expertise or prior knowledge of robot swarms often struggle to anticipate or correctly perceive fragmentation events (cf. section 2). One way to help such operators better perceive and anticipate swarm fragmentation *and* react accordingly to prevent it could be to make certain aspects of the behavioural model of swarm robots explicit, as proposed by Kolling et al. [25]. Regardless of the specific operators task, conveying low-level information about the robots' behaviour that the operators may not perceive, could improve their understanding of the swarm dynamics resulting from the robot interactions [16, 25], and therefore help them anticipate and prevent fragmentation. Such low-level information is usually communicated at the robot level, through the use of localised visualisations (see section 2.1).

The main objective of this study is to assess how localised visualisations of robot behaviour influence a human operator's ability to prevent swarm fragmentation. We frame the robot's behaviour using the Perception–Decision–Action (PDA) loop [56], a common model in multi-agent systems. Within this framework, we designed three visualisations that make explicit what each robot perceives, how it processes that information to make decisions, and what actions it takes as a result. These visualisations were evaluated in a user study with 51 participants. The evaluation is conducted in an immersive Virtual Reality environment with simulated robots. Immersive VR is used not only for experimental control but also as a proxy for future AR deployment in real-world scenarios. While evaluated in immersive Virtual Reality, these visualisations are designed with co-located Augmented Reality overlays in mind, enabling direct mapping of robot behaviour to their physical location in future real-world deployments. In addition, four swarm behaviours previously used in the literature [13, 15, 49] are implemented to present various situations of fragmentation to the participants. Fragmentation situations are triggered by changes in the model parameters, which disrupt the swarm's self-organisation dynamics. All these conditions are evaluated to assess their impact on users' ability to 1) anticipate fragmentation and 2) choose the appropriate control to prevent it.

The study presented in this document is structured as follows. First, a review of the literature on swarm fragmentation and visualisations is presented in section 2, followed by section 3 which focuses on the information and visualisations related to the PDA behaviour loop of the robots. The experimental protocol is then detailed in section 4 and the results are presented in section 5. Finally, an interpretation of the results and a discussion about limitations and perspectives are provided in section 6, followed by a conclusion in section 7.

2 RELATED WORK

Fragmentation [21, 31, 35] is a well-known phenomenon in swarm systems, and can pose significant risks when it occurs in contexts where cohesion is critical to mission success. When a human operator is in charge, preventing inappropriate fragmentation becomes essential. As a result, understanding swarm behaviour and anticipating fragmentation remain key challenges for effective human-swarm interaction. Previous Human-Swarm Interaction (HSI) studies [5, 6, 15] have been conducted to determine how humans perceive swarm fragmentation. Capiola et al. [5, 6] reported that humans are not proficient enough to accurately detect swarm fragmentation or the proportion of the swarm that is fragmented in a flocking simulation. Hénard et al. [15] showed that fragmentation is generally well-perceived (87% accuracy) except during expansion behaviour, likely due to uniform spatial distribution that masks connectivity loss. They highlighted the role of Gestalt principles, especially proximity and common fate, in fragmentation perception. As an alternative, visually augmenting the robots themselves—through co-located visualisations that display relevant information directly where the robots are—may help improve the operator's understanding of the

swarm. Augmented reality (AR) provides a natural way to overlay such information in the operator's field of view, supporting more intuitive interpretations of swarm behaviour. These visual augmentations aim to enhance swarm transparency and improve the operator's awareness at three levels: what is happening, why it is happening, and what might happen next [7, 39]. This section reviews existing visualisation strategies for robot swarms, with a focus on localised displays and the use of Augmented Reality in Human–Swarm Interaction. It also discusses how visual information is managed to support user performance and presents common experimental setups used for evaluation.

2.1 Robot swarm visualisations

To the best of our knowledge, recent taxonomies [45, 48] classifying visualisations used on robots seem difficult to apply to swarms, as they are mainly aimed at a single robot and do not take into account the characteristics of multi-robot systems, such as the interaction between robots or the spatial distribution specific to these systems, which require special attention. As visualisations applied to swarms of robots are characterised according to various aspects, we propose some criteria for classifying them. The first criterion is the visualisation level. As the situation may need to consider the swarm as a whole, an aggregation of subsets or individuals clearly discriminated, the visualisations can characterise several levels within the swarm, such as the robot level [12, 34], the group level [38] or the swarm level [40, 50]. The second criterion is the visualisation frequency. As a lot of robots may be augmented, the user's field of vision could be saturated, as could their mental load. Consequently, it may be appropriate for some visualisations to be displayed conditionally, while others are displayed systematically [40, 50]. The third criterion is the type of information conveyed by the visualisation. Visualisations may display robots' capabilities such as robots' perception area [11, 26, 44]; robots' internal decision process, such as the current state of the robots when controlled by a Finite State Machine [28]; interactions and relationships within the swarm such as groups [2, 8, 34], links between robots [2, 40, 52], the force of influence between robots [32] and swarm's leaders [50]; information about the current action of the swarm such as robots' direction [12, 26, 34, 50], swarm's direction [40, 50] or robot's trajectories [3, 29]; information about the spatial distribution of the swarm such as swarm's centroid [50], swarm densities [40] or swarm outline [40, 41, 50]; state prediction such as robot's future position [51]; environmental influence such as force fields [17, 18, 23]. A last criterion could be to identify if the visualisation aims to add new information [12, 41], convey information already present in a different way [40, 50], or remove information from the swarm, by hiding some of the robots for example [50].

2.2 Localised visualisations and Augmented Reality in Human–Swarm Interaction

An increasing number of studies explore the use of localised and co-located visualisations in HSI, particularly through Augmented Reality (AR). Unlike global views, AR allows information to be anchored directly on or around the robots, improving users' spatial understanding and reducing context-switching [12, 28]. These visualisations have been used for debugging [28], trajectory analysis [12], and group behaviour interpretation [2], often revealing internal robot states or inter-robot interactions. Beyond analysis, AR interfaces also enable natural interaction with robot swarms through gesture, gaze, or parameter tuning, with varying levels of control granularity [33, 55]. Multi-user systems, such as swARm [20], have begun to explore collaborative control with head-mounted displays. Broader reviews highlight the promise of AR in robotics, while underlining challenges in visual clutter, latency, and task relevance [27, 30]. While promising, few studies have systematically assessed how localised AR cues affect user perception and decision-making. Our study contributes to this gap by evaluating visualisations grounded in the Perception–Decision–Action (PDA) loop, with a focus on their utility for anticipating and preventing swarm fragmentation in immersive settings.

2.3 Managing information complexity in swarm visualisation

Various strategies have been proposed to optimise visualisations in Human–Swarm Interaction, not only based on the type of information users need, but also on how these visualisations affect performance. A key challenge is to balance informational completeness with cognitive load, especially when monitoring large swarms.

One common approach is to reduce the amount of visual information to avoid overwhelming the user. This can be done by selectively hiding certain robots [50] or replacing detailed views with global summaries [38, 40]. However, the effectiveness of this strategy varies: it may degrade performance [40], have no measurable impact [50], or even improve decision-making and situational awareness (SA) [38], depending on the task. The potential trade-off lies in removing critical details that support user understanding of swarm behaviour [50].

Another approach is to augment the display with additional information—such as anticipated states or inferred relationships between agents. This strategy can enhance performance [41, 51], provided the information is relevant and well-integrated. Otherwise, it risks increasing cognitive load and hindering interpretation [40].

In both cases, the usefulness of a visualisation depends on what is shown, how it is shown, and when it is shown. Effective design therefore requires identifying which information can be removed, transformed, or added to support user performance without causing overload—an aspect that is especially critical when deploying localised visual cues in AR.

2.4 Experimental set-up

Studying interactions between humans and robot swarms can be challenging, as it often requires deploying a large number of physical robots, which entails technical complexity and logistical constraints. Several approaches have been proposed to facilitate such research. A commonly used solution is to visualise the swarm on a 2D screen from a bird's eye view [15, 29, 40, 50]. This method is practical and scalable, but may not capture key aspects of user experience, such as embodiment or spatial perception. When factors such as the user's point of view, the realism of the swarm, and its interaction with the environment are central to the study, physical robots may be used [2, 10, 34, 41]. While offering ecological validity, such setups are often harder to reproduce and more difficult to control experimentally. An alternative approach is to use Virtual Reality (VR) [4, 46], which allows participants to be immersed in a realistic simulation of the environment and the swarm [24]. VR enables precise control over experimental variables, facilitates reproducibility, and offers the ability to monitor the task from the participant's perspective, ensuring consistent conditions across trials.

Summary and design guidelines. From these previous studies, we can identify several relevant principles for designing visualisations that help users detect and anticipate swarm fragmentation. First, spatial *proximity* and *common fate*, as grounded in Gestalt theory—appear to play a critical role in human perception of swarm structure. Second, the approach in AR—naturally being to augment the display with localised additional information on real robots—may increase cognitive load, but it may also convey information that is essential to the operator's awareness [16]. Third, visualisations should enhance perceptual salience of connectivity or motion divergence to support fragmentation anticipation. However, most prior works focus on 2D or non-immersive environments, with limited consideration for real-time AR constraints such as visual clutter or field-of-view limitations. In this work, we build on these findings and explore three visualisations aligned with distinct stages of the perception-decision-action loop, aiming to test their usability and effectiveness in an immersive, simulated AR context.

3 ANALYSIS OF PERCEPTION-DECISION-ACTION VISUALISATIONS

To support the operator's understanding of swarm dynamics, it can be useful to provide visualisations that represent different aspects of the robots' behaviour. In this work, we rely on the Perception–Decision–Action (PDA) loop [56], a widely used framework in the

field of multi-agent systems, to structure and present these visualisation strategies for robot swarms.

As collective behaviour at the swarm level (macro) emerges from individual robot interactions (micro), visualising information at the level of perception, decision, and action may help users bridge these two levels of abstraction. This idea echoes the recommendations of Kolling et al. [25], who advocate for exposing key aspects of the swarm's internal model to improve user supervision. Such visualisations can also support the development of a coherent mental model [22], enabling users to describe, explain, and predict the swarm's behaviour more effectively.

Furthermore, making these internal processes visible contributes to increasing system transparency [39], which is essential for effective human–swarm interaction. In particular, aligning low-level behavioural cues with high-level outcomes may improve users' ability to anticipate critical events—such as fragmentation—and to select appropriate interventions, as discussed in Section 2.3.

The PDA loop [56] describes an agent's behaviour in three stages: perception (environmental input), decision (internal processing), and action (output behaviour). We apply this model to swarm robots to map internal states to external visual cues. When the swarm's behaviour leads to its fragmentation, the decisive information for anticipating its appearance may be found in one of these three aspects. Using visualisations to convey this information to users may be a solution to prevent fragmentation.

3.1 Robot perception

Among the information perceived by the robots in the environment, some represent other robots of the swarm. As long as the robots in the swarm perceive each other, they can adapt their behaviour according to their neighbours. For users, robot connectivity cannot be observed explicitly, and it is possible that humans do not perceive the connectivity of the swarm, as suggested in [15].

Connectivity can be visualised at the local robot level by displaying the robots' perception fields [11, 26, 44] or the links between robots [2, 40]. But it can also be visualised at the global swarm level by displaying the convex hull of connected robots [40, 41, 50] or the groups by colouring them using distinct colours [8, 34]. However, these visualisations at the global level only provide new information on the swarm connectivity *after fragmentation*. For example, group colouring shows a single colour when all robots are connected, and two or more colours only after fragmentation occurs. To help anticipate fragmentation, the visualisation must provide information on the connectivity of the swarm before fragmentation, as with these local-level visualisations. In both cases, visualisations mostly use the grouping principle named *Uniformed Connectedness* from the *Gestalt Theory* [47], naturally making humans identifying groups in a simple glance by using links [40], colours [34] or convex hull [41] for example.

3.2 Robot decision

The decision of the robot is an internal process that takes into account the robot's perception of its environment and itself, to determine its next action. In human-swarm interaction, users usually follow a training approach [14] to build a mental model of this decision-making process. Providing users with real-time information about this process could help them understand and anticipate the dynamics of the swarm, thus improving the effectiveness of their training, or even replacing it. However, there are few examples in the literature of swarms transmitting what the robots decide. An example can be found in [28] where robots are coloured according to the current state of their Finite State Machine. For example, for a robot coming dangerously close to another robot, showing that it is currently avoiding robots or an obstacle can indicate to users that its cohesion parameter is not too high and that they should instead reduce its repulsion parameter to avoid collision. In swarms where robots are driven by virtual forces [15], their behaviour does not rely on discrete states, but rather emerges from the continuous combination of multiple interacting forces. Visualising which force currently dominates a robot's movement can provide insights into the swarm's underlying dynamics [25], which may help operators anticipate and

Visualisation	PDA Component	Displayed Information
Li (Links)	Perception	Explicit neighbour links showing robot-to-robot connectivity and the strength of these connections.
Fo (Forces)	Decision	Dominant interaction forces, derived from local agent rules (e.g., attraction, repulsion)
Di (Directions)	Action	Robot movement direction

Table 1: Mapping between the three proposed visualisations and the components of the Perception–Decision–Action loop.

prevent fragmentation. This approach is not limited to force-based models—it can also be applied to other control architectures. For example, in systems based on fuzzy logic—a method that uses degrees of truth rather than binary rules to model behaviour—highlighting which behavioural rules are currently active can help users understand what drives the robots’ actions at any given time.

3.3 Robot action

The actions that can be carried out by swarm robots are generally limited. One example is the deposition of virtual pheromones, which can be made visible to the operator [19]. However, the most common and fundamental action in most swarm systems is movement.

In this context, visualising the “action” component of the PDA loop typically involves displaying information related to robot motion. This can be done at the individual level—by showing robot orientation [10, 12, 26, 32, 34, 50], past trajectories [3, 29], or predicted future positions [51]—or at the swarm level, such as by displaying the average direction or speed of the swarm [40]. Because swarm motion emerges from the collective movement of individual robots, visual cues at one level (individual or group) may be adapted or scaled to the other. While some of these cues—like movement direction—may be naturally perceived by human observers [54], others might require visual enhancement to support interpretation and decision-making.

4 EXPERIMENTAL PROTOCOL

Building on the PDA-based visualisations described in the previous section, the present experiment investigates whether such visual cues can help human operators better manage swarm cohesion. Specifically, it addresses two main research questions: **(Q1)** Does providing information through PDA-based visualisations improve the operator’s ability to anticipate fragmentation? **(Q2)** Does this information help the operator select appropriate control strategies to prevent fragmentation?

To standardise stimuli across participants, we used pre-recorded clips covering four swarm behaviours: flocking, swarming, expansion, and densification. Using these behaviours will allow us to answer the following secondary question: **(Q3)** What is the impact of the behaviour on the ability of humans to prevent swarm fragmentation?

To answer these questions, we set up a within-subjects experimental design based on two separate tasks, using Virtual Reality (VR) to immerse the users in the operator role [46]. In the first task, participants have to determine whether the presented simulated swarms will fragment. In the second task, participants have to determine which of the two proposed controls will prevent the fragmentation. Three visualisations (see section 4.2) are compared during the two tasks to assess their impact on users’ performance.

4.1 Swarm behaviour model

For this study, the implementation detailed in [15] based on virtual forces is reproduced. This model provides flocking, swarming, expansion and densification behaviours, by using six behavioural rules and modifying the weight of each rule. Three of these rules respectively represent cohesion, separation and alignment from the Reynolds’ *Boids* model [37]. The other three are random movement (allows the robot




Fig. 2: Screenshot of the simulation showing virtual Mona robots, one of which has its perception zone displayed (330° field of vision, 0.6 m perception distance)

to move independently and randomly), friction (reduces the agent’s speed over time) and overlapping avoidance (prevents robots from overlapping). The sum of the forces produced by these six rules gives the robot’s acceleration. The acceleration is then used to update the robot’s velocity (direction and speed). Cohesion, separation, alignment and overlapping avoidance are rules that depend on the robot’s perceived neighbourhood, *i.e.* other robots in its perception zone. This zone is circular and is influenced by two parameters, its radius and a blind spot located at the rear of the robot (cf. Fig. 2).

4.2 Evaluated visualisations

To represent the three components of the swarm model’s PDA loop, three distinct visualisations were implemented, following the analysis presented in section 3. These visualisations are applied at the robot level, as opposed to the swarm level. This design choice is particularly important for the *perception* aspect: as discussed in section 3, swarm-level representations of perception do not provide useful information prior to the onset of fragmentation, whereas robot-level visualisations can reveal early cues related to local sensing and interaction. In addition, as this experiment aims to check if the aspects shown are usually missing for humans, the visualisations are displayed systematically. Based on these choices, we proposed three visualisations and compared them with each other and with a baseline (Ba) condition in which no additional visual information is provided. In total, four experimental conditions were tested. The three proposed visualisations respectively convey: (1) the current state of connectivity between robots (*perception*), (2) the dominant virtual force influencing each robot (*decision*), and (3) the movement direction of each robot (*action*). To clarify how each visualisation aligns with the components of the PDA loop, we summarise their purpose and mapping in Tab. 1.

4.2.1 Visualisations of links (Li)

Two types of information are conveyed by this first visualisation of the robots’ perception aspect: what are the connected robots, and how far each connection is from the break. Two robots are connected if at least one of them perceives the other. The chosen visualisation represent links between robots as solid grey lines. In addition, the line’s thickness is modified according to the distance between the robots, until it disappears when the connection is broken (see Fig. 1.b). Instead of just indicating which robots are interacting (*e.g.* [2, 40]), qualifying the distance between two robots might provide relevant information for users, particularly to anticipate fragmentation.

4.2.2 Visualisation of dominant forces (Fo)

In the behavioural model, the *dominant force* is the virtual force with the greatest intensity at time t , thus having the strongest influence on a robot’s acceleration. This excludes the random component, which is unaffected by neighbourhood interactions and fluctuates independently. To visualise the decision process, a coloured directional arrow is displayed above each robot (see Fig. 1.c), encoding the dominant force’s nature, direction, and relative intensity. The arrow measures 0.26 m in length, 0.05 m in height, with a 0.15 m head and 0.05 m

tail. Direction indicates acceleration; colour encodes force type and strength: red (attraction), blue (repulsion), green (alignment), and white (no dominant force). Colour intensity reflects alignment between force and acceleration: fully saturated at 0° , fading to white by 90° , beyond which the force is not considered dominant. By indicating the prevailing influence in each robot’s local decision-making, this visualisation aims to support users in understanding the swarm’s internal dynamics and choosing the most effective interventions to prevent fragmentation.

4.2.3 Visualisation of directions (Di)

Information about the robots’ actions consist in the current direction of each robot, as implemented in previous works (e.g. [32,34,50]). The direction is obtained using the velocity vector of each robot, calculated by the swarm model (see section 4.1). These directions are displayed using a pink directional arrow, whose dimensions are the same as in the previous visualisation (see Fig. 1.d). As humans are naturally able to perceive the movement of robots, this visualisation should not provide users with any additional information to help them anticipate or prevent fragmentation.

4.3 Swarm recordings

We recorded 24 stimuli clips (50 fps, 6–7 s each; mean 6.75 s) showing a simulated swarm in motion. Half of them lead to fragmentation, half do not. Robots move on a 2D plane (max speed 0.2 m/s, median 0.1 m/s), with a 330° field of view and either 0.6 or 1.0 m perception range. Clips were created by trimming longer simulations: fragmented clips end 0.5 s before disconnection, and matched non-fragmented clips were selected to resemble them in density, spatial distribution and motion. Fragmentation was programmatically verified using a connectivity graph: a clip is marked as fragmented if the graph is disconnected. Each clip is associated with one of four behaviours (densification, expansion, flocking, swarming), with six clips per behaviour (three with fragmentation, three without). Clips are viewed in VR using a Meta Quest 2 headset in two separate tasks. Guided randomisation controls their distribution. To prevent recognition across conditions, clips are mirrored horizontally (relative to the participant) in half the presentations. The same version of a clip is never shown consecutively. The 6–7 s duration ensures participants have time to perceive swarm dynamics without inducing fatigue.

4.4 Environment

The environment (see Fig. 1) is a virtual office measuring 11x13 m, containing a 7x7 m arena where the robots can move around, bounded by borders 0.1 m thick and 0.2 m high. Depending on the value of the robots’ perception radius used in the experiment, the perception area of each robot covers either 2.3 or 6.4% of the arena. The participant stands and is located outside the arena in a static position, 1 m from the edge of the arena, and halfway along an edge. The 3D models representing the robots are Mona [1] robots measuring 0.08 m in diameter and 0.03 m in height (see Fig. 2). All the elements create and are subject to shadows.

4.5 Procedure

First, participants are given information about the experiment and then sign the consent form if they wish to continue. They then fill in the demographic form. The concept of swarm fragmentation is introduced to the participants, along with a video showing a swarm fragmenting. Different colours for distinguishing the different sub-swarms are used in the video to highlight fragmentation. No explicit information on the swarm model is given to the participant during the experiment. Next, the experimenter explains the first task to the participants and makes sure that they understand it. The participants then read a document briefly explaining the visualisations and providing the required information to use them, before putting on the headset and starting the first task. Once the first task has been completed, the second task is explained to the participants along with an explanatory video. Once again, the participants put on the headset and complete the second task. Once the last task has been completed, participants complete a feedback questionnaire (see section 4.5.3).

4.5.1 First task: fragmentation anticipation

The first task aims to study the participants’ ability to anticipate swarm fragmentation depending on the different visualisations proposed. For this task, participants must watch each of the clips (see section 4.3) presented to them and answer the question “Will the swarm fragment?”. Participants watch each clip in its entirety before being able to respond. When the clip ends, the swarm disappears and the response interface appears in the environment. They then have as much time as they need to give their answer, by clicking on the answer of their choice on the virtual interface. Participants are informed before the start of the task that if fragmentation has to occur, it must happen soon after the end of the clip. During this task, the participant’s responses and response times are recorded. Performance in this task was measured objectively as the proportion of correct answers, based on whether the presented swarm ultimately fragmented or remained cohesive in the simulation.

4.5.2 Second task: appropriate control

The second task aims to evaluate whether participants can correctly select the appropriate control—between two proposed options—to prevent swarm fragmentation. It also seeks to assess the influence of visualisations on their decision-making. For this task, only the 12 clips that eventually lead to fragmentation are used (see Section 4.3). Importantly, fragmentation is not yet visible in the clips, as each ends just before it occurs. Each of these 12 clips is shown once per visualisation condition, resulting in 48 decision-making trials per participant. As in the first task, participants watch each clip in its entirety and are then prompted with the question: “*What changes should have been made to the swarm to prevent it from fragmenting?*” The interface appears immediately after the clip ends. To avoid learning effects, participants do not receive feedback on their choices and cannot observe the consequences of their selected control. The two proposed responses for each clip are predetermined based on the swarm’s state. One proposal effectively prevents fragmentation; the other does not. We chose parameter changes [25] as the means of control because they are directly linked to the swarm model, providing the low-level manipulation required for an operator role. This low-level access means participants need a functional understanding of swarm dynamics to select the appropriate parameters. To define the two control options, we used our simulator to resume the scenario from the clip and observe the outcomes of different controls. Each clip’s random seed is fixed, ensuring reproducibility of the swarm’s behaviour. Although the controls are context-specific, they are drawn from a limited set of four possible interventions: *Increase cohesion*, *Reduce cohesion*, *Increase alignment*, and *Reduce alignment*. For statistical consistency, only one correct control is provided per trial. This choice is not always the most intuitive one, making the task deliberately challenging to better reveal the impact of the visualisations. Prior to the task, participants watch a single explanatory video demonstrating how each control affects swarm behaviour. This video contains no fragmentation and is shown only once, with no further explanation provided by the experimenter. Consequently, participants are given only macroscopic-level cues about swarm behaviour and must infer the micro-level robot interactions themselves. Task 2 performance is quantified as the percentage of trials in which participants selected the control that, in simulation, successfully prevented fragmentation.

4.5.3 Qualitative data

At the end of the experiment, the participants fill in a questionnaire designed to gather qualitative results. First, participants are asked to give their perception of their performance on each of the two tasks. In addition, they must indicate which visualisation they found most helpful for each task. The participants are then asked to explain what they understood about the effect of each of the controls from the second task. Finally, they can express themselves freely about the experiment and are free to add any additional comments.

4.6 Hypotheses

Based on the findings of the literature and the analysis provided in section 3, we formulate the following hypotheses regarding the expected effects of the visualisations evaluated in this study. Each visualisation

Table 2: Summary of Independent and Dependent Variables for both tasks.

Independent Variables (common to both tasks)			
participants	51	(random variable)	
visualisation type	4	baseline, links, forces, directions	(within-subject)
swarm behaviour	4	densification, expansion, flocking, swarming	(within-subject)
fragmentation status	2	present / absent	(within-subject, task 1 only)
clip repetition	4	Each clip shown once per visualisation	
Dependent Variables			
<i>Task 1 – Fragmentation Anticipation</i>			
anticipation accuracy	Correct answer to “Will the swarm fragment?”		
response time	Time to respond after clip ends		
<i>Task 2 – Control Selection</i>			
control accuracy	Correct identification of the control preventing fragmentation		
response time	Time to respond after clip ends		

targets a specific aspect of the PDA loop, and their impact is assessed on two tasks: the ability to anticipate the appearance of fragmentation, and the ability to choose the appropriate control to prevent it.

For the **connectivity visualisation**, which corresponds to the perception aspect, we expect that making robot connectivity visible will help participants detect early signs of fragmentation. We thus hypothesise that it will improve anticipation scores (**H1-1**), but will not significantly influence control selection (**H1-2**).

For the **dominant force visualisation**, which reflects the decision-making process, we anticipate that it will not help participants anticipate fragmentation (**H2-1**), but will improve their ability to choose an appropriate control action to prevent it (**H2-2**).

Finally, for the **direction visualisation**, which illustrates the action aspect of the PDA loop, we do not expect a significant impact on either anticipation (**H3-1**) or control selection scores (**H3-2**), as directional motion is generally well perceived by humans without additional visual aid.

4.7 Preparations

The ethics committee of the authors’ institution has approved this experimental protocol. The study was conducted in a quiet room exclusively dedicated to the study for its duration. A total of 51 participants with no prior expertise in swarms took part in the experiment, of which 69% were male and 31% were female. Among them, 35% took part in experiments involving swarms of robots. The participants ranged in age from 18 to 56 (median of 25). All participants had a correct or corrected vision. Collected data were anonymised.

5 RESULTS

The impact of two factors, visualisations (Q1 & Q2) and swarm behaviours (Q3), was tested in the experiment on the ability to anticipate fragmentation (first task) and to choose the appropriate control (second task). Non-parametric tests were used, as the experiment uses a within-subject pattern and a Shapiro test has confirmed that data do not follow a normal distribution. The Friedman test followed by a Nemenyi post-hoc were used to check differences in correct answers and response times among visualisation. Non-significant results are not reported in this document. To assess the impact of the different swarm behaviours (Q3), all the tests are done in the baseline condition, which represents a quarter of the participants’ data.

5.1 Results of the anticipation task

While analysing the results of the first task, six participants were removed from the analysis as outliers using the interquartile range method. Consequently, the following tests were performed on the results obtained from 45 participants.

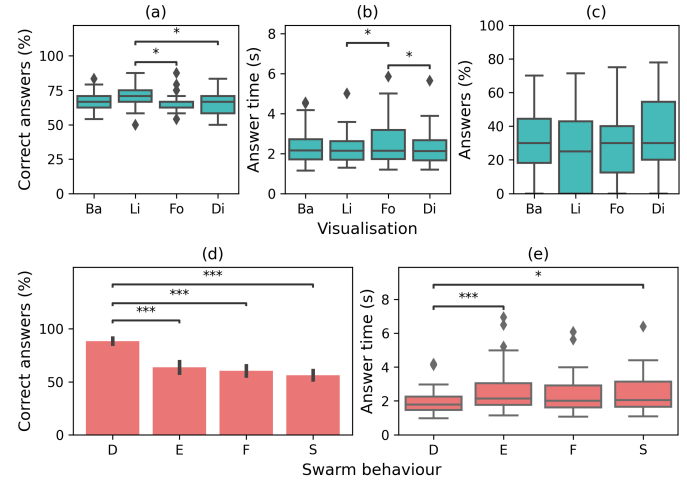


Fig. 3: Diagrams showing results for the first task, mainly by showing the medians and quartiles: (a) correct answer rate, (b) answer time, and (c) among errors, rate of ‘will fragment’ responses, according to the visualisation displayed; (d) correct answers rate with mean values and standard deviations, and (e) answer time, according to the behaviour displayed by the swarm. (D) Densification (E) Expansion (F) Flocking (S) Swarming

5.1.1 Impact of visualisations

When analysing the effect of the visualisations on participants’ accuracy in anticipating fragmentation, the Friedman test shows statistically significant differences ($Q = 13.67, p = 0.003$). The Nemenyi post-hoc test indicates that **the links visualisation** significantly improved fragmentation anticipation compared to both the dominant force ($p = 0.011$) and direction ($p = 0.015$) visualisations. Figure 3.a shows the results.

Then, when examining the answer time of participants in anticipating fragmentation, the Friedman test shows statistically significant differences ($Q = 10.63, p = 0.01$). The Nemenyi post-hoc test shows that **the dominant force visualisation** was found to significantly increase the answer time than either visualisation displaying **links** ($p = 0.028$) or **directions** ($p = 0.022$). Figure 3.b shows the obtained results. Finally, looking at the wrong answers from participants in anticipating fragmentation, the Friedman test cannot be used because the number of wrong answers is different for each visualisation. Instead, the non-parametric alternative of the One-way ANOVA, the Kruskal-Wallis test, is used. The Kruskal-Wallis test shows no statistically significant differences in wrong answers type ($Q = 3.68, p = 0.29$). Figure 3.c shows the obtained results.

5.1.2 Impact of swarm behaviour

When analysing the effect of the swarm’s behaviours on participants’ accuracy in anticipating fragmentation, the Friedman test shows statistically significant differences ($Q = 66.42, p < 0.001$). The Nemenyi post-hoc test shows that fragmentation from the densification category was found to be significantly easier to anticipate than either expansion ($p < 0.001$), flocking ($p < 0.001$) or swarming ($p < 0.001$). Figure 3.d shows the obtained results. Then, looking at the answer time from participants in anticipating fragmentation, the Friedman test shows statistically significant differences ($Q = 17.98, p < 0.001$). The Nemenyi post-hoc test shows that densification behaviour was found to significantly reduce the answer time than either expansion ($p < 0.001$) or swarming ($p = 0.013$). Figure 3.e shows the obtained results.

5.2 Results of the appropriate control task

While analysing the results of the second task, 2 participants were removed from the analysis as outliers using the interquartile range method. Consequently, the following tests were performed on the results obtained from 49 participants.




Fig. 4: Diagrams showing results for the second task, mainly by showing the medians and quartiles: (a) correct answers rate, and (b) answer time, according to the visualisation displayed; (c) correct answers rate with mean values and standard deviations, (d) answer time, according to the behaviour displayed by the swarm. (D) Densification (E) Expansion (F) Flocking (S) Swarming

5.2.1 Impact of visualisations

When analysing the effect of visualisations on participants’ accuracy in choosing the appropriate control, the Friedman test shows no statistically significant differences ($Q = 5.80, p = 0.12$). Figure 4.a shows the obtained results. Then, looking at the answer time from participants, the Friedman test shows statistically significant differences ($Q = 9.17, p = 0.027$). The Nemenyi post-hoc test shows that **the links visualisation** was found to significantly increase the answer time in choosing the appropriate control than the baseline ($p = 0.015$). Figure 4.b shows the obtained results.

5.2.2 Impact of swarm behaviour

When analysing the effect of the swarm’s behaviours on participants’ accuracy in choosing the appropriate control, the Friedman test shows statistically significant differences ($Q = 12.81, p = 0.005$). The Nemenyi post-hoc test shows that it is significantly easier to choose the right interaction to prevent fragmentation from the densification category than swarming ($p = 0.028$). Figure 4.c shows the obtained results. Then, looking at the answer time from participants, the Friedman test shows statistically significant differences ($Q = 22.34, p < 0.001$). The Nemenyi post-hoc test shows that densification behaviour was found to significantly reduce the answer time in choosing the control than either flocking ($p = 0.007$) or swarming ($p < 0.001$). Moreover, the expansion behaviour was found also to have a significant difference compared to the swarming behaviour ($p = 0.007$). Figure 4.d shows the obtained results.

5.3 Qualitative results

The qualitative results were obtained from the participants using a form. The results presented in this section are based on the responses of 51 participants. For each task, participants were asked which visualisation they considered the most helpful in completing the task. They were also asked to give their opinion of their performance on each task on a five-point scale ranging from “1-Not at all” to “5-Perfectly”. Finally, they were asked to explain the effect of the controls in the second task in their own words.

Feedback on the first task Regarding the distribution of the preferred visualisations for the first task (Fig. 5.a), the **links visualisation** has the highest proportion (73%), compared to displaying direction (12%), dominant force (14%) or the baseline (2%). Concerning participants’ opinions of their performance in the first task, the mean value is 2.78 (median = 3, see Fig. 5.c).

Feedback on the second task Regarding the distribution of the preferred visualisations for the second task (Fig. 5.b), the **dominant force visualisation** has the highest proportion (47%), compared to displaying links (25%), direction (24%) or the baseline (4%). shows the obtained results. Concerning participants’ opinions of their performance in the second task, the mean value is 3.03 (median = 3, see Fig. 5.d).




Fig. 5: Diagrams showing qualitative results: (a) visualisation most useful for the first task, and (b) visualisation most useful for the second task, according to participants; (c) participants’ opinion of their success in the first task, on a scale of 1-Not at all to 5-Perfectly; (d) participants’ opinion of their success in the second task, on a scale of 1-Not at all to 5-Perfectly.

Controls understanding Two principles were used in the second task controls, cohesion and alignment. Participants were asked to explain the effect of increasing or decreasing these two principles independently of each other. In their answer, it was checked whether the participants correctly explained the effect of the principle on the swarm and whether they noticed that the effect applies to the robots’ neighbourhood. Of the explanations given by participants for the cohesion principle, 84% mentioned an effect on robots moving closer together or further apart, and 39% mentioned the involvement of the robot’s neighbours. Concerning the alignment principle, 64% of participants mentioned an effect on the coordination of robot movement, and 27% mentioned the involvement of the robot’s neighbours. Then, it was checked whether the participants highlighted the effect of control on fragmentation. 8% of participants mentioned using the controls for specific fragmentation situations: increasing cohesion to avoid fragmentation when robots spread out too much, decreasing cohesion when several groups form, increasing alignment when two groups go in distinct directions or when robots move randomly, and decreasing alignment when they are too far apart and keeping their current direction prevents them from moving closer together. Finally, the number of participants with an explanation that was undoubtedly incorrect was checked. 2% of the participants incorrectly explained the effect of the cohesion principle, as they mentioned the synchronisation of the robots’ movement and 10% of the participants incorrectly explained the effect of the alignment principle, as they mentioned changes in distance between robots.

6 DISCUSSION

Out of the six tested hypotheses (three per task), only one was clearly supported: the link-based visualisation (Li) significantly improved fragmentation anticipation. This perception-level cue outperformed the direction- and force-based visualisations, which showed limited or no advantage over the baseline. Participants also considered the link visualisation the most useful, suggesting that low-level connectivity information supports anticipation more effectively than decision- or action-level cues. However, this effect was not consistent across tasks: the link visualisation did not significantly improve performance in the control task. These mixed results indicate that while perception-level cues can aid anticipation, decision-level and action-level visualisations may require further refinement, simplification, or contextual adaptation to become effective. Regarding swarm behaviours, fragmentation caused by densification was both easier to anticipate and to control than fragmentation caused by expansion, flocking or swarming. This may reflect stronger visual salience or more interpretable collective dynamics in densification scenarios.

6.1 Impact of visualisations

Links. This visualisation was expected to be more helpful to the users than the others for anticipating fragmentation (H1-1), as it was designed to show just how close two agents are to breaking up, but finally shows no significant difference to the baseline. One participant commented

that displaying the links between the robots gives an impression of cohesion which may be contradictory to the movement of the robots, which seems to indicate future fragmentation. This is consistent with the fact that the display of links relies on the notion of *uniform connectedness* from *Gestalt theory* [47] to identify robots that are part of the same group. However, no effect was found in the results to corroborate this hypothesis (see Fig. 3.c). This suggests that this visualisation only conveys the connectivity of the swarm, but not the dynamics of their interaction. A conditional visualisation showing only the links of interest or a visualisation focusing on the dynamics of the interaction between the robots might be more appropriate in this context. Finally, the same reasoning could be applied in the context of the second task, as this visualisation shows no improvement, as expected (H1-2).

Dominant forces. This visualisation aimed to expose internal decision influences of the swarm model [25], to help users select appropriate control actions (H2-2). However, results show no benefit over the baseline, and five participants reported that it was disturbing. Displaying the dominant force for every robot often led to clutter, especially under swarming dynamics. Arrows varied in direction and colour—even within the same neighbourhood—and changed constantly, making it difficult to identify relevant zones and analyse the swarm in time. Excess information increases cognitive load, potentially leading users to disengage from the task [57]. The challenge lies in designing visualisations that are both faithful to the swarm model and cognitively efficient [36]. While our study used VR, the goal is AR supervision—where co-located, persistent cues must remain legible and lightweight. The poor performance of this visualisation underscores the risk of overwhelming users with low-level, dynamic data. A contributing factor may be a mismatch between the visualisation and the task demands. Although participants seemed to form mental models of the swarm dynamics [22], the second task required not just understanding, but prediction—anticipating how a control would affect swarm behaviour. The visualisation may have supported comprehension but failed to aid forward inference. Several participants described the arrows as dense and difficult to interpret. In AR, this would be worsened by reduced field of view and real-time constraints. Future work should explore adaptive strategies: filtering or aggregating arrows, selectively displaying salient cues, or shifting to higher-level representations aligned with human attentional focus.

Directions. In general, displaying directions is useful for representing movement when it is absent or imperceptible [53]. Still, it was not expected (H3-1 and H3-2) in this context to provide the users with any more information than they already perceive, since humans are naturally able to perceive moving elements [54]. The results obtained support this, showing no significant difference compared to the baseline condition. Despite the results obtained, the transmission of information on the action aspect of the swarm other than the directions of the robots (e.g., the relative movements of the robots with respect to the swarm) could give users the information they need to understand the dynamics of the swarm.

These results highlight how the PDA framework can help structure the design and evaluation of swarm visualisations. While the perception-oriented visual cue (connectivity) appears more immediately useful for fragmentation anticipation, the decision and action layers (dominant force and direction) present more nuanced design challenges. This supports the idea that the PDA loop is not only a modelling tool, but also a useful lens to guide the mapping of swarm internals to human-perceivable cues in AR. The lack of effect for the force and direction visualisations may be due to a mismatch between what is represented and what participants need to decide. For instance, force vectors—though theoretically informative—may have been cognitively demanding to interpret within the short clip duration, especially under time pressure. Similarly, motion direction arrows may not sufficiently highlight deviations that predict fragmentation.

6.2 Swarm behaviours

Among the four behaviours compared in this study, densification is significantly easier to anticipate (see Fig. 3.d), and the one with the best score for choosing the right interaction, even though there was

no significant difference (see Fig. 4.c). A difference is noticeable in comparing these results with those of Hénard et al. [15]. In their study, participants were free to anticipate or not their response to identify fragmentation. In the end, the authors noted that participants were more likely to wait until the last moment to respond when the swarm adopted swarming behaviour. In this study, the participants were forced to respond after viewing the clip. Swarming did not turn out to be significantly more difficult to anticipate than the other behaviours. It was the densification behaviour that stood out, being easier than the others to prevent. This result is consistent with the results from [15], as fragmentation in densification behaviour proved easier to identify. The authors suggest that the notion of *common Fate* from *Gestalt theory* [47] plays an important role concerning densification behaviour, as the identification of one or more centres of gravity seems to be a crucial clue for the identification of fragmentation in this behaviour. For the other three behaviours, the absence of visual evidence of fragmentation (e.g. group identification through the detection of distinct directions or variations in distance, respectively based on *common fate* and *proximity* grouping principles) makes the task more difficult. The expansion behaviour shows a synchronised and homogeneous movement, the swarming behaviour shows a disordered movement that is difficult to predict, and the flocking behaviour shows an ordered and synchronised movement that shows few or no anomalies. As variations in distance between robots occur without immediately causing fragmentation, the densification behaviour is probably the only one to give the users enough time to anticipate fragmentation. Finally, despite the difference in the user's point of view compared to [15], and the use of an immersive environment, the results are consistent.

6.3 Limitations and future work

This study has several limitations that also outline practical takeaways for future work and deployment. First, participants were non-experts in swarm supervision, as recruiting trained operators is challenging. Results may differ with professional users, so operational studies—for instance, with drone fleet coordinators or underwater robot pilots—are needed to validate the findings. Our second task aimed to assess control selection, but its design (limited options, no feedback) may have restricted differences between visualisations. Future iterations should mirror real scenarios, where operators receive feedback and adjust commands dynamically. Although we used virtual reality to simulate immersive AR, true AR deployment raises extra challenges: limited field of view, motion parallax, display instability, and environment lighting may all affect usability. Evaluating the system on AR headsets in field conditions (e.g., for search-and-rescue teams or agricultural monitoring) is a crucial next step. Scalability and spatiality are also concerns. Our study involved moderately sized, mostly planar swarms, but in real-world scenarios—especially with aerial or underwater drones—robots may operate in 3D formations. Local cues may clutter at scale or fail to convey altitude changes, so visualisations should integrate layering, filtering, or clustering to remain interpretable. Finally, our results suggest that visual cues should complement—not duplicate—perceptible swarm behaviour. Redundant or overly dense displays risk increasing cognitive load. Instead, adaptive, task-specific, and perceptually efficient encodings—for example, emphasizing connectivity when fragmentation risk rises, or abstracting forces into simplified “corridors” of movement—will better support robust human-swarm interaction.

7 CONCLUSION

This paper investigated the impact of three localised visualisations—each corresponding to a component of the Perception–Decision–Action (PDA) loop—on human ability to anticipate and prevent swarm fragmentation. Through an immersive VR study involving 51 participants, we evaluated whether conveying information about robot connectivity (perception), dominant behavioural influence (decision), or movement direction (action) could support operators across two tasks: anticipating fragmentation and selecting an appropriate control to avoid it. Results reveal that visualising inter-robot connectivity improves anticipation performance, both quantitatively and qualitatively, confirming the utility of perception-level cues for

identifying early signs of fragmentation. In contrast, displaying dominant forces or directions did not significantly enhance user performance, and none of the three visualisations consistently improved control selection. These limitations may stem from visual overload or redundancy with what humans naturally perceive, particularly in the case of direction cues. Densification scenarios were easier for participants to interpret and act upon than expansion, flocking, or swarming, which remain visually ambiguous and cognitively demanding.

This study illustrates how the PDA loop can be used as a conceptual structure for designing swarm visualisations. However, its operational benefit varies depending on the component visualised and the task context. In our setting, perception-level visualisation (Li) showed the most promise for aiding anticipation, though its advantage did not extend to all tasks, nor did it consistently outperform the baseline. Taken together, these findings underline the importance of prioritising connectivity cues for early fragmentation detection, simplifying or avoiding redundant force and direction displays that risk overloading users, and ensuring that visualisations are designed with adaptability and scalability in mind for dense, dynamic swarms. These insights provide actionable guidance for the design of co-located visual feedback in Augmented Reality (AR), where visual clutter and attentional load are critical concerns. Future work should explore adaptive strategies tailored to swarm dynamics, integrate predictive or global representations, and test the approach under real AR conditions. By aligning visual feedback with operators' mental models and attentional capacity, AR-based systems could better support effective human oversight in high-density or rapidly evolving swarm scenarios.

SUPPLEMENTAL MATERIALS

Supplemental materials are available at <https://osf.io/49gny>. They notably include three videos showing details of the experiment, and the document explaining the visualisations, given to participants at the beginning of the experiment. In addition, the code and simulation parameters used in the study are available upon request from the authors, to ensure full reproducibility of the experiment.

ACKNOWLEDGMENTS

This work is part of the ANR ARTUISIS (<https://siia.univ-brest.fr/artuisis/>) project, which received funding under the reference ANR-21-CE33-0006.

REFERENCES

- [1] F. Arvin, J. Espinosa, B. Bird, A. West, S. Watson, and B. Lennox. Mona: an affordable open-source mobile robot for education and research. *Journal of Intelligent & Robotic Systems*, 94:761–775, 2019. doi: 10.1007/s10846-018-0866-9 5
- [2] S. Batra, J. Klingner, and N. Correll. Augmented reality for human–swarm interaction in a swarm-robotic chemistry simulation. *Artificial Life and Robotics*, 27(2):407–415, 2022. doi: 10.1007/s10015-022-00763-w 2, 3, 4
- [3] D. Biediger, A. Mahadev, and A. T. Becker. Investigating the survivability of drone swarms with flocking and swarming flight patterns using virtual reality. In *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, pp. 1718–1723, 2019. doi: 10.1109/COASE.2019.8843173 2, 4
- [4] B. Capelli, V. Villani, C. Secchi, and L. Sabatini. Understanding multi-robot systems: on the concept of legibility. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7355–7361, 2019. doi: 10.1109/IROS40897.2019.8968023 3
- [5] A. Capiola, I. a. Hamdan, E. L. Fox, J. B. Lyons, K. Sycara, and M. Lewis. “is something amiss?” investigating individuals’ competence in estimating swarm degradation. *Theoretical Issues in Ergonomics Science*, 23(5):562–587, 2022. doi: 10.1080/1463922X.2021.1983887 2
- [6] A. Capiola, D. Johnson, I. a. Hamdan, J. B. Lyons, and E. L. Fox. Detecting swarm degradation: Measuring human and machine performance. In J. Y. C. Chen and G. Fragomeni, eds., *Virtual, Augmented and Mixed Reality*, pp. 325–343. Springer Nature Switzerland, Cham, 2023. 2
- [7] J. Y. Chen, K. Procci, M. Boyce, J. Wright, A. Garcia, and M. Barnes. Situation awareness-based agent transparency. *US Army Research Laboratory*, (April):1–29, 2014. 2
- [8] M. Chen, P. Zhang, Z. Wu, and X. Chen. A multichannel human-swarm robot interaction system in augmented reality. *Virtual Reality & Intelligent Hardware*, 2(6):518–533, 2020. doi: 10.1016/j.vrih.2020.05.006 2, 3
- [9] S. Chen, D. Yin, and Y. Niu. A survey of robot swarms’ relative localization method. *Sensors*, 22(12):4424, 2022. doi: 10.3390/s22124424 1
- [10] M. Daily, Y. Cho, K. Martin, and D. Payton. World embedded interfaces for human-robot interaction. In *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*, pp. 6 pp.–, 2003. doi: 10.1109/HICSS.2003.1174285 3, 4
- [11] A. M. R. Denniss, T. E. Gorochoowski, and S. Hauert. Augmented reality for the engineering of collective behaviours in microsystems. In *2019 International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS)*, pp. 1–6, 2019. doi: 10.1109/MARSS.2019.8860907 2, 3
- [12] F. Ghiringhelli, J. Guzzi, G. A. Di Caro, V. Caglioti, L. M. Gambardella, and A. Giusti. Interactive augmented reality for understanding and analyzing multi-robot systems. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1195–1201, 2014. doi: 10.1109/IROS.2014.6942709 2, 4
- [13] J. Harvey, K. E. Merrick, and H. A. Abbass. Assessing human judgment of computationally generated swarming behavior. *Frontiers in Robotics and AI*, 5:13, 2018. doi: 10.3389/frobt.2018.00013 2
- [14] J. D. Hasbach, T. E. F. Witte, and M. Bennewitz. On the importance of adaptive operator training in human-swarm interaction. In R. A. Sottolare and J. Schwarz, eds., *Adaptive Instructional Systems*, pp. 311–329. Springer International Publishing, Cham, 2020. 3
- [15] A. Hénard, E. Peillard, J. Rivière, S. Kubicki, and G. Coppin. Human perception of swarm fragmentation. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction, HRI ’24*, 9 pages, p. 250–258. Association for Computing Machinery, New York, NY, USA, 2024. doi: 10.1145/3610977.3634961 2, 3, 4, 8
- [16] A. J. Hepworth, D. P. Baxter, A. Hussein, K. J. Yaxley, E. Debie, and H. A. Abbass. Human-swarm-teaming transparency and trust architecture. *IEEE/CAA Journal of Automatica Sinica*, 8(7):1281–1295, 2020. 2, 3
- [17] T. Hiraki, S. Fukushima, Y. Kawahara, and T. Naemura. Phygital field: An integrated field with physical robots and digital images using projection-based localization and control method. *SICE Journal of Control, Measurement, and System Integration*, 11(4):302–311, 2018. doi: 10.9746/jcmsi.11.302 2
- [18] T. Hiraki, S. Fukushima, Y. Kawahara, and T. Naemura. Navigatorch: Projection-based robot control interface using high-speed handheld projector. In *SIGGRAPH Asia 2019 Emerging Technologies, SA ’19*, 3 pages, p. 31–33. Association for Computing Machinery, New York, NY, USA, 2019. doi: 10.1145/3355049.3360538 2
- [19] E. R. Hunt, S. Jones, and S. Hauert. Testing the limits of pheromone stigmergy in high-density robot swarms. *Royal Society open science*, 6(11):190225, 2019. doi: 10.1098/rsos.190225 4
- [20] H. James. *swARm: Augmented Reality Swarm Control with Multiple Head Mounted Display Operators*. Doctoral dissertation, Worcester Polytechnic Institute, 2021. 2
- [21] L. G. Jeston-Fenton, S. Abpeikar, and K. Kasmarik. Visualisation of swarm metrics on a handheld device for human-swarm interaction. In *2022 26th International Conference Information Visualisation (IV)*, pp. 149–154, 2022. doi: 10.1109/IV56949.2022.00032 2
- [22] N. A. Jones, H. Ross, T. Lynam, P. Perez, and A. Leitch. Mental models: an interdisciplinary synthesis of theory and methods. *Ecology and society*, 16(1), 2011. 3, 8
- [23] J. Kato, D. Sakamoto, M. Inami, and T. Igarashi. Multi-touch interface for controlling multiple mobile robots. In *CHI ’09 Extended Abstracts on Human Factors in Computing Systems, CHI EA ’09*, 6 pages, p. 3443–3448. Association for Computing Machinery, New York, NY, USA, 2009. doi: 10.1145/1520340.1520500 2
- [24] H. Khodr, U. Ramage, K. Kim, A. Guneysoz Ozgur, B. Bruno, and P. Dillenbourg. Being part of the swarm: Experiencing human-swarm interaction with vr and tangible robots. In *Proceedings of the 2020 ACM Symposium on Spatial User Interaction, SUI ’20*, article no. 23, 2 pages. Association for Computing Machinery, New York, NY, USA, 2020. doi: 10.1145/3385959.3422695 3
- [25] A. Kolling, P. Walker, N. Chakraborty, K. Sycara, and M. Lewis. Human interaction with robot swarms: A survey. *IEEE Transactions on Human-Machine Systems*, 46(1):9–26, 2016. doi: 10.1109/THMS.2015.2480801 2, 3, 5, 8

- [26] K. Libby, L. Zhong, and N. Van Stralen. Diagnosing robotic swarms 2 (dr. swarm2). Retrieved December, 13:2020, 2020. 2, 3, 4
- [27] Z. Makhataeva and H. A. Varol. Augmented reality for robotics: A review. *Robotics*, 9(2):21, 2020. 2
- [28] A. G. Millard, R. Redpath, A. M. Jewers, C. Arndt, R. Joyce, J. A. Hilder, L. J. McDaid, and D. M. Halliday. Ardebug: An augmented reality tool for analysing and debugging swarm robotic systems. *Frontiers in Robotics and AI*, 5, 2018. doi: 10.3389/frobt.2018.00087 2, 3
- [29] D. Miner and N. Kasch. Swarmvis: a tool for visualizing swarm systems. *UMBC Computer Science*, 636, 2008. 2, 3, 4
- [30] D. Naser Addin and B. Ozell. Design and test of an adaptive augmented reality interface to manage systems to assist critical missions, 2021. arXiv preprint arXiv:2103.06676. 2
- [31] R. Olfati-Saber. Flocking for multi-agent dynamic systems: algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, 2006. doi: 10.1109/TAC.2005.864190 1, 2
- [32] E. Pajorová, L. Hluchý, and M. Masár. 3d visualization of the simulation research results of new method for cooperative of a group of mobile agents for swarm behavior simulation. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ASONAM '13, 5 pages, p. 1275–1279. Association for Computing Machinery, New York, NY, USA, 2013. doi: 10.1145/2492517.2500291 2, 4, 5
- [33] J. Patel and C. Pinciroli. Improving human performance using mixed granularity of control in multi-human multi-robot interaction. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pp. 1135–1142. IEEE, 2020. 2
- [34] G. Podevijn, R. O'Grady, and M. Dorigo. Self-organised feedback in human swarm interaction. In *Proceedings of the workshop on robot feedback in human-robot interaction: how to make a robot readable for a human interaction partner (Ro-Man 2012)*, 2012. 2, 3, 4, 5
- [35] G. Punzo, J. Simo, D. J. Bennet, and M. Macdonald. Characteristics of swarms on the edge of fragmentation. *Phys. Rev. E*, 89:032903, 5 pages, Mar 2014. doi: 10.1103/PhysRevE.89.032903 2
- [36] C. L. Remondino, B. Stabellini, P. Tamborini, et al. Visualizing complex systems. In *Visualizing Complex Systems*, pp. 17–31. Politecnico di Torino, 2019. 8
- [37] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '87, 10 pages, p. 25–34. Association for Computing Machinery, New York, NY, USA, 1987. doi: 10.1145/37401.37406 1, 4
- [38] K. A. Roundtree, J. R. Cody, J. Leaf, H. Onan Demirel, and J. A. Adams. Visualization design for human-collective teams. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 63, pp. 417–421. SAGE Publications Sage CA: Los Angeles, CA, 2019. 2, 3
- [39] K. A. Roundtree, M. A. Goodrich, and J. A. Adams. Transparency: Transitioning from human-machine systems to human-swarm systems. *Journal of Cognitive Engineering and Decision Making*, 13(3):171–195, 2019. 2, 3
- [40] K. A. Roundtree, M. D. Manning, and J. A. Adams. Analysis of human-swarm visualizations. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 62(1):287–291, 2018. doi: 10.1177/1541931218621066 2, 3, 4
- [41] S. O. Sachidanandam, S. Honarvar, and Y. Diaz-Mercado. Effectiveness of augmented reality for human swarm interactions. In *2022 International Conference on Robotics and Automation (ICRA)*, pp. 11258–11264, 2022. doi: 10.1109/ICRA46639.2022.9812080 2, 3
- [42] E. Şahin. Swarm robotics: From sources of inspiration to domains of application. In E. Şahin and W. M. Spears, eds., *Swarm Robotics*, pp. 10–20. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. 1
- [43] J. Scholtz. Theory and evaluation of human robot interactions. In *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*, pp. 10–pp. IEEE, 2003. doi: 10.1109/HICSS.2003.1174284 1
- [44] M. Sugimoto, T. Fujita, H. Mi, and A. Krzywinski. Robotable2: A novel programming environment using physical robots on a tabletop platform. In *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology*, ACE '11, article no. 10, 8 pages. Association for Computing Machinery, New York, NY, USA, 2011. doi: 10.1145/2071423.2071436 2, 3
- [45] R. Suzuki, A. Karim, T. Xia, H. Hedayati, and N. Marquardt. Augmented reality and robotics: A survey and taxonomy for ar-enhanced human-robot interaction and robotic interfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, CHI '22, article no. 553, 33 pages. Association for Computing Machinery, New York, NY, USA, 2022. doi: 10.1145/3491102.3517719 2
- [46] V. Villani, B. Capelli, and L. Sabatini. Use of virtual reality for the evaluation of human-robot interaction systems in complex scenarios. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 422–427, 2018. doi: 10.1109/ROMAN.2018.8525738 3, 4
- [47] J. Wagemans, J. H. Elder, M. Kubovy, S. E. Palmer, M. A. Peterson, M. Singh, and R. von der Heydt. A century of gestalt psychology in visual perception: I. perceptual grouping and figure-ground organization. *Psychological bulletin*, 138(6):1172, 2012. doi: 10.1037/a0029333 3, 8
- [48] M. Walker, T. Phung, T. Chakraborti, T. Williams, and D. Szafir. Virtual, augmented, and mixed reality for human-robot interaction: A survey and virtual design element taxonomy. *J. Hum.-Robot Interact.*, 12(4), article no. 43, 39 pages, jul 2023. doi: 10.1145/3597623 2
- [49] P. Walker, M. Lewis, and K. Sycara. Characterizing human perception of emergent swarm behaviors. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 002436–002441. IEEE, 2016. doi: 10.1109/SMC.2016.7844604 2
- [50] P. Walker, M. Lewis, and K. Sycara. The effect of display type on operator prediction of future swarm states. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 002521–002526, 2016. doi: 10.1109/SMC.2016.7844619 2, 3, 4, 5
- [51] P. Walker, S. Nunnally, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara. Neglect benevolence in human control of swarms in the presence of latency. In *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 3009–3014. IEEE, 2012. doi: 10.1109/ICSMC.2012.6378253 2, 3, 4
- [52] P. M. Walker. *Improving operator recognition and prediction of emergent swarm behaviors*. PhD thesis, University of Pittsburgh, 2017. 2
- [53] W. H. Warren, A. Blackwell, K. Kurtz, N. G. Hatsopoulos, and M. Kalish. On the sufficiency of the velocity field for perception of heading. *Biological cybernetics*, 65:311–320, 1991. doi: 10.1007/BF00216964 8
- [54] S. N. Watamaniuk and A. Duchon. The human visual system averages speed information. *Vision research*, 32(5):931–941, 1992. 4, 8
- [55] B. Williamson, E. Taranta, Y. Moolenaar, and J. LaViola. Command and control of a large scale swarm using natural human interfaces. *Field Robotics*, 3:301–322, 2023. 2
- [56] M. Wooldridge and N. R. Jennings. Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995. doi: 10.1017/S0269888900008122 2, 3
- [57] V. Yoghoudjian, Y. Yang, T. Dwyer, L. Lawrence, M. Wybrow, and K. Marriott. Scalability of network visualisation from a cognitive load perspective. *IEEE transactions on visualization and computer graphics*, 27(2):1677–1687, 2020. 8