



HAL
open science

Evaluation of augmented visualisations to prevent the fragmentation of robot swarm

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

► **To cite this version:**

Aymeric Hénard, Jérémy Rivière, Etienne Peillard, Sébastien Kubicki, Gilles Coppin. Evaluation of augmented visualisations to prevent the fragmentation of robot swarm. 2024. hal-04541882

HAL Id: hal-04541882

<https://hal.science/hal-04541882v1>

Preprint submitted on 11 Apr 2024

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.

Evaluation of augmented visualisations to prevent the fragmentation of robot swarm

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

Abstract—Swarm fragmentation is a breakdown of communication and coordination among robots in the swarm. Detrimental for the self-organisation of the swarm, it can have multiple causes such as the robots’ limited capabilities, obstacles in the environment, wrong model parameter settings or unfitting control by an operator. Transmitting information about the swarm using visualisations to human operators could allow them to prevent these fragmentations. In this paper, we propose three augmented visualisations respectively based on the perception, decision process, and action of each robot. We evaluate their benefit on the ability of humans to anticipate and prevent fragmentation arising from four swarm behaviours selected amongst expansion, densification, flocking and swarming. The effect of each of these behaviours is also evaluated in the same context. The experimental protocol consists of two tasks, in which participants have to 1) anticipate the apparition of fragmentation, and 2) choose the appropriate command to prevent fragmentation from occurring. The results show that visualisation showing links between connected robots outperforms visualisation showing the direction of the robots or the force dominating their movement in anticipating fragmentation, but no visualisation proves to be more useful in preventing fragmentation. In addition, they show that humans anticipate fragmentations arising from densification behaviour more accurately than from the other three behaviours. Our study reveals elements that will be useful in designing future visualisation for human-swarm interaction.

Index Terms—Human-Swarm Interaction, Robot swarm, Visualisation, Control, User study, Virtual Reality

I. INTRODUCTION

Swarms of robots are autonomous systems composed of many robots with low capabilities [1] that exhibit self-organised collective behaviours, such as flocking [2], which require the coordination of the robots that compose them. However, the technical limitations of the robots only allow them to communicate over limited distances [3]. Some robots may no longer be able to interact, and the swarm may split up. This phenomenon is known as swarm fragmentation [4]. 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. The users can adopt different roles when interacting with the swarm [5],

and in particular the *operator* role, which implies low-level control and short-term actions, making it possible to influence the self-organisation of the swarm. If the self-organisation of the swarm is the cause of the fragmentation, influencing it can enable the users to prevent fragmentation. In the context of preventing fragmentation, displaying the appropriate visualisations may enable the operators to anticipate or at least perceive swarm fragmentation [6] and react accordingly. Using an Augmented Reality (AR) device could improve the operators’ understanding of the swarm dynamics by displaying information directly on the robots [7].

The main objective of this study is to evaluate how augmented visualisations based on the robots’ behaviour affect the human capability to prevent swarm fragmentation. Making certain aspects of the behavioural model of swarm robots explicit could convey the relevant information to users to facilitate the understanding of swarm dynamics, as proposed by [8], and therefore help them anticipate and prevent fragmentation. These aspects of the model can be characterised by the Perception-Decision-Action (PDA) loop [9] of multi-agent systems. In this context, three visualisations representing what each robot perceives, what drives their decision process and what is the decided action are evaluated in a user study involving 51 participants. In addition, four swarm behaviours previously used in the literature [6], [10], [11] are implemented to present various situations of fragmentation to the participants. These situations of fragmentation are triggered by changes in the parameters of the model, causing a disturbance in the swarm’s self-organisation. All these conditions are evaluated to determine how they affect the ability of users to anticipate the appearance of fragmentation, and also whether they affect the ability of users to choose the appropriate control to prevent their appearance. The obtained results suggest that it is preferable to focus the choice of visualisation on the perception aspect of the robots’ PDA loop rather than on the action or decision aspect to help users anticipate fragmentation, and that densification behaviour is easier to anticipate than the other three. 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 II, followed by section III which focuses on the information and visualisations related to the PDA behaviour loop of the robots. The experimental protocol is then detailed in section IV and the results are presented in section V. Finally, an interpretation of the results and a discussion about limitations and perspectives are provided in section VI, followed by a conclusion in section VII.

A. Hénard and J. Rivière are with Univ Brest, CNRS, Lab-STICC, Brest, France

É. Peillard and G.Coppin are with IMT Atlantique, CNRS, Lab-STICC, Brest, France

S. Kubicki is with ENIB, CNRS, Lab-STICC, Brest, France.

II. RELATED WORK

Fragmentation [4], [12], [13] is a well-known phenomenon in swarms and raises important risks when it is not appropriate for the situation and the context. When humans are in charge of the swarm, they should prevent fragmentation that could be detrimental to the accomplishment of the mission. Therefore, the question of understanding the behaviour of a swarm and detecting or even preventing any fragmentation remains crucial. Previous Human-Swarm Interaction studies [6], [14], [15] have been conducted to determine how humans perceive swarm fragmentation. In [14], [15], the authors 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. In [6], the authors evaluate the human ability to perceive fragmentation among four swarm behaviours: flocking, swarming, expansion, and densification. Overall, the results show that swarm fragmentation is well-perceived for most behaviours, except for expansion, which is more challenging. The authors suggest that expansion behaviour is difficult to assess because the visual appearance of the swarm is the same in a fragmentation or cohesion situation, which both exhibit uniform behaviour and spatial distribution, making the connectivity of the swarm hard to perceive. They also show that participants have difficulty anticipating the appearance of swarm fragmentation, possibly because the elements used to assess the presence of fragmentation appear clearly once it has already occurred. To help the users prevent fragmentation, different solutions can be explored. Firstly, dedicated metrics could be developed that predict the onset of fragmentation [12], [13], warning users of potential risk. However, using such a solution is unlikely to explain to the users why fragmentation occurs and what can be done to avoid it. As swarm behaviour and fragmentation result from the individual behaviour of robots and their local interactions, perceiving their presence does not mean that users understand why they appear.

An alternative solution could be to visually augment the robots, by adding appropriate localised information. This section presents an overview of the visualisations used in the literature related to robot swarms, and the experimental setups used to evaluate them.

A. Robot swarm visualisations

To the best of our knowledge, recent taxonomies [16], [17] 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 [18], [19], the group level or the swarm level [20],

[21].

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 [20], [21], and the whole cognitive load must be seriously considered.

The third criterion is the type of information conveyed by the visualisation. Among the existing visualisations applied to swarms of robots, we can find different types of information. Visualisations may display:

- robots' capabilities such as robots' perception area [22]–[24]
- robots' internal decision process, such as the current state of the robots when controlled by a Finite State Machine [25]
- interactions and relationships within the swarm such as groups [18], [26], [27], links between robots [21], [27], [28], the force of influence between robots [29] and swarm's leaders [20]
- information about the current action of the swarm such as robots' direction [18]–[20], [24], swarm's direction [20], [21] or robot's trajectories [30], [31]
- information about the spatial distribution of the swarm such as swarm's centroid [20], swarm densities [21] or swarm outline [7], [20], [21]
- state prediction such as robot's future position [32]
- environmental influence such as force fields [33]–[35]

A last criterion could be to identify if the visualisation aims to add new information [7], [19], convey information already present in a different way [20], [21], or remove information from the swarm, by hiding some of the robots for example [20].

B. Visualisation choice according to context

Previous works on Human-Swarm Interaction evaluate the use of different visualisations in different contexts. As some swarms can be composed of many robots, some works tested to reduce the amount of information conveyed to users using different strategies [20], [21]. The results of these studies show that hiding swarm members by displaying only some of them [20] or displaying specific aspects of the swarm using a global representation [21] is detrimental to users if it is not appropriate to the current task. Indeed, using a global representation can be equivalent to observing swarm members if the global representation conveys appropriate information to users, such as synthesised information about swarm spatiality in a task about identifying swarm spatiality [20]. By default, it seems preferable to display the members of the swarm, which is suitable for most contexts, rather than a global representation of the swarm which would only be suitable for a specific context, as the latter has not proved more useful [20], [21]. In addition, the use of a visualisation showing the members of the swarm enables additional information to be displayed on the swarm, which may normally be imperceptible to users, which could improve their performance in their task [7].

C. Experimental set-up

Deploying an experimental set-up to study interactions between humans and robot swarms can be complicated, as it often requires the use of many robots. Different solutions exist to carry out research on swarms. The solution widely used is to display the swarm members on a 2D screen with a bird's eye view [6], [20], [21], [30]. But in situations where the user's point of view, the realism of the swarm and its environment are important experimental factors, a solution is to use real robots [7], [18], [27], [36]. An alternative is to use a Virtual Reality (VR) environment [37], [38] to immerse the users into a reproduction of real-world conditions. The use of VR also allows experimenters to control the experimental variables, ensure the reproducibility of the experiment and follow the task from the participant's eyes to ensure that everything is going correctly.

III. ANALYSIS OF PERCEPTION-DECISION-ACTION VISUALISATIONS

The PDA loop [9] is a way of representing the functioning of the robots' behaviour. This concept was developed in the field of Multi-Agent Systems and allows the cyclical behaviour of an agent to be compartmentalised into 3 aspects. For each robot, this loop encompasses the robot's perception, its decision and its action. The perception aspect represents all the information obtained from the environment, whether raw or interpreted, using sensors or communications. The decision aspect represents the process to determine the next action of the robot, using the information perceived and the robot's internal information. Finally, the action aspect represents the execution of the action decided by the robot. Once the robot has acted, its environment changes, triggering the PDA loop again.

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.

A. 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. However, for users, the connectivity of robots is not something that can be observed explicitly, and it is possible that humans do not perceive the connectivity of the swarm, as suggested in [6]. Connectivity can be visualised at the local robot level by displaying the robots' perception fields [22]–[24] or the links between robots [21], [27]. But it can also be visualised at the global swarm level by displaying the convex hull of connected robots [7], [20], [21] or the groups by colouring them using distinct colours [18], [26]. 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 case, visualisations mostly use the grouping principle named *Uniformed Connectedness* from the *Gestalt Theory* [39], naturally making humans identifying groups in a simple glance by using links [21], colours [18] or convex hull [7] for example.

B. 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. Users usually follow a training approach [40] 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 [25] 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 it should instead reduce its repulsion parameter to avoid collision. In the case of a swarm whose robots use virtual forces [6], the behaviour of the robots is not made up of discrete states, as different forces drive the behaviour of the robots. Indicating which force dominates the robot's movement may inform the users about the dynamics of the swarm [8], which could be useful for preventing fragmentation. And it also works for other types of implementation. For example, by indicating the active rules of robots controlled by fuzzy logic, we can inform users about what causes the robots to act.

C. Robot action

The actions that can be carried out by swarm robots are limited. An example of these actions is the deposit of pheromones, which can be visualised by displaying the pheromones [41]. But the most common action performed by swarm robots is to move. In this case, showing the action aspect of the PDA loop corresponds to displaying information related to their movement. Visualising the movement of the swarm can be done at the robots' level by showing the direction of the robots, which is frequently used in the literature. [18]–[20], [24], [36], [42], their past trajectories [30], [31], or their potential future position [32]. But it can also be done at the swarm level, by showing the average direction of the swarm [21] or its average speed [21]. In addition, the visualisations displayed at the robot level can be applied at the group or swarm level and vice versa, as the movement of a swarm is the composition of the movement of its robots. These visualisations convey various information to users, some of which could already be naturally perceived by humans, such as robots' directions given that they are naturally able to see movement as well as variations in movement within a set of moving entities [43].

IV. EXPERIMENTAL PROTOCOL

The experiment conducted in this study aims to determine if the proposed visualisations (see section IV-B) can help prevent swarm fragmentation. More precisely two research questions are addressed here:

- (Q1) Does providing information through visualisations based on the PDA loop improve human anticipation of fragmentation?
- (Q2) Does providing information through visualisations based on the PDA loop help humans choose the appropriate control to prevent fragmentation?

To ensure the swarms in motion presented are the same for all participants, we chose to pre-record clips from the swarm simulation covering different categories of behaviour such as 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 role of operator [37]. The first task is a binary discrimination 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 fragmentation from occurring. Three visualisations (see section IV-B) are compared during the two tasks to assess their impact on users' performance.

A. Swarm behaviour model

In this study, the implementation detailed in [6] based on virtual forces is used. This model provides flocking, swarming, expansion and densification behaviour, 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 [2]. The other three are random movement (allows the robot 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 speed (direction and intensity). Cohesion, separation, alignment and overlapping avoidance are rules that depend on the perceived neighbourhood of the robot applying them. This neighbourhood corresponds to the robots present in the robot 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. In this study, two robots are considered to be connected if at least one of them perceives the other.

B. Evaluated visualisations

To represent the three aspects of the swarm model's PDA loop, three visualisations are implemented based on the analysis from section III. The visualisations are displayed at the

robot level, notably because visualisations at the swarm level for the perception aspect do not give any information before the appearance of fragmentation, as detailed in section III-A. 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, three visualisations are chosen and are compared in the experiment with each other and with a baseline, giving a total of four conditions. The chosen visualisations are the following:

1) *Links*: To represent the perception aspect, the links between the connected robots are displayed. As well as showing the links between robots, the line's thickness is modified according to the distance between the robots, until it disappears when the connection is broken (see Figure 1.b), to show users how far the robots are from the break. Instead of just indicating which robots are interacting, qualifying the distance between two robots might provide relevant information for users, particularly to anticipate fragmentation.

2) *Dominant force*: To represent the decision aspect, the dominating forces of the robots are displayed, showing between attraction, repulsion, alignment or the lack of domination, using a coloured directional arrow above each robot: red, blue, green and white respectively (see Figure 1.c). The arrow has the following dimensions: 0.26 m long, 0.05 m high, 0.15 m head width, 0.05 m tail width. It points in the direction of the robot's acceleration. Moreover, the stronger the force, the more intense the colour. Here are the steps followed to obtain the dominant force for each robot:

- 1) From the forces composing the acceleration, the one with the greatest magnitude is identified.
- 2) The final colour depends on this force.
- 3) To compute the intensity of the colour, the angle between the acceleration and the force is calculated.
 - If the angle obtained is 0° , then the intensity of the colour is maximum.
 - As the angle increases, the colour approaches white, until it becomes white at 90° or more.

If the angle is greater than 90° , the force is not considered dominating the movement.

3) *Direction*: To represent the action aspect, the robots' directions are displayed. In this experiment, the direction of each robot is displayed above them with the same arrow, as for the display of the direction (see Figure 1.d). The direction is obtained using the speed vector of each robot, calculated by the swarm model (see section IV-A).

4) *Baseline*: Finally, displaying only the robots without adding any additional information allows us to test the participants' anticipation capabilities and provide a comparison basis for assessing the impact of the three other visualisations. This visualisation therefore represents the baseline for this study, as robots are displayed in all visualisations (see Figure 1).

C. Swarm recordings

The stimuli presented to the participants are recordings of a swarm in motion. A total of 24 different clips were recorded at 60 frames per second, varying in length from 6 to 7 seconds (6.75 seconds on average). Among them, 12 lead to a




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 direction using pink directional arrows




Fig. 2. Environment of the experiment, (a) Axis of symmetry used to reverse robot positions (b) Participant position during both tasks.

fragmentation of the swarm and 12 do not, for which the model used is defined in section IV-A. The simulated robots navigate on a plane at a maximum speed of 0.2 m/s (median: 0.1 m/s), and have a 330° field of vision and 0.6 or 1.0 m perception distance, depending on the clip. To obtain the clips with and without fragmentation, longer versions are recorded, before being cut to the desired length. Clips containing fragmentation end 0.5 seconds before fragmentation. For those containing no fragmentation as a reference point, they are cut so that the swarm dynamics correspond to those with fragmentation. For each recording, fragmentation is identified using a function based on the robots’ perception, by representing robots as graph nodes and interactions between robots as links. If the resulting graph is disconnected and contains two or more components, it means the swarm is fragmented. The clips are also classified according to four behaviours that can be obtained with the model used: densification, expansion, flocking and swarming. Each behaviour is represented by six clips, half of which lead to swarm fragmentation. These clips are used in two separate tasks, and viewed by participants in VR (see Figure 1) using a Meta Quest 2 head-mounted display. During these tasks, guided randomisation ensures that clips and visualisations are distributed correctly. To prevent participants from recognising the clips despite the different visualisations, a right-left symmetry is used (in relation to the participant’s position), which does not cause any change in perspective or distance. This symmetry is used on two of the four presentations of the same clip. The same version is never shown consecutively.

D. Environment

The environment (see Figure 2) 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. The participant stands and is located outside the arena in a fixed position, 1 m from the edge of the arena, and standing halfway

along an edge. The 3D models representing the robots are Mona [44] robots measuring 0.08 m in diameter and 0.03 m in height. All the elements create shadows and are subject to them.

E. 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 video¹ 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 (see section IV-F) 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 (see section IV-G). 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 IV-H)

F. 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 IV-C) 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.

G. Second task: Appropriate Control

The second task aims to study whether participants can choose the right control, between two proposals, to prevent fragmentation. It also aims to show the impact of visualisations on participants' choices. Consequently, only the 12 clips leading to swarm fragmentation (see section IV-C) are used for this task. Fragmentation is still not visible in the clips, as they end before it. As each clip is viewed once for each category of visualisation, participants have to choose the appropriate control in 48 situations. In the same way, as for the first task, participants watch a clip in its entirety and then answer the question: "What changes should have been made to the swarm to prevent it from fragmenting?" on the interface that

appears once the clip has finished. The proposed responses to prevent fragmentation of the swarm at each clip are defined beforehand, as they depend on the state of the swarm. One of the two proposals prevents fragmentation, the other does not. We chose parameter changes [8] as a means of control because it is directly linked to the model, and therefore offers the low-level control required for the operator role. In addition, because of its low-level control over the model, users must have an adequate understanding of the dynamics of the swarm to adjust the parameters correctly to obtain the desired result. To define the two proposals, we used our simulator to switch from the clip to the simulation, enabling us to check which control prevents the appearance of fragmentation and which does not. To ensure that the proposed controls lead to the desired result on the swarm, the seed of the random generator is saved for each clip and the execution of the system is controlled to obtain the same simulation each time. Even if the proposals are specific to the clips, they are selected from among four control alternatives: "Increase cohesion", "reduce cohesion", "increase alignment" and "reduce alignment". However, one clip in the swarming category has a different proposal from the others, because a combination of cohesion and alignment was needed to prevent fragmentation. The proposal for this clip is: "increase cohesion and alignment" or "reduce cohesion and increase alignment". For the other clips, it is sometimes possible to prevent fragmentation using different approaches. As only one control to prevent fragmentation can be proposed to participants, the control chosen is not necessarily the most obvious one to prevent fragmentation. This choice makes the test deliberately difficult to better show the impact of the visualisations. These controls are explained beforehand using a video², showing the effect of the control on a swarm. There is no fragmentation in this video. In addition, the participants only watch the video once and are not given any further explanation by the experimenter. Consequently, participants were given only explicit information on the macroscopic level of the swarm. As they were not given any information about how the swarm model works, they had to deduce by themselves how the robots interact with each other at the micro-level.

H. 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.

I. Hypotheses

Based on the findings of the literature and the analysis provided in section III, assumptions are made about the results

¹REMOVED FOR ANONYMISATION

²REMOVED FOR ANONYMISATION

that can be expected from the visualisations evaluated in this study. The main hypotheses, null or alternative, concerning visualisations are presented below:

- H1:1-3 (Alternative hypothesis) The rate of correct answers obtained by the visualisation showing links for the anticipation task is different from that obtained by:
 - the baseline (H1:1)
 - the visualisation showing direction (H1:2)
 - the visualisation showing dominant force (H1:3)
- H1:4-6 (Alternative hypothesis) The rate of correct answers obtained by the visualisation showing the dominant force for the appropriate control task is different from that obtained by:
 - the baseline (H1:4)
 - the visualisation showing links (H1:5)
 - the visualisation showing direction (H1:6)
- H0:7-8 (Null hypothesis) The rate of correct answers obtained by the visualisation showing the direction is equal to that obtained by the baseline for the:
 - anticipation task (H0:7)
 - appropriate control task (H0:8)

These hypotheses represent our expectations prior to the evaluation. However, all the conditions of the experiment are evaluated in this study and not just these hypotheses.

J. 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.

V. RESULTS

Two factors were tested in the experiment: the impact of

- the different visualisations (Q1 & Q2),
- and of the different swarm behaviours (Q3),

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.

A. Results of the anticipation task

While analysing the results of the first task, 6 participants were removed from the analysis, as they were considered to

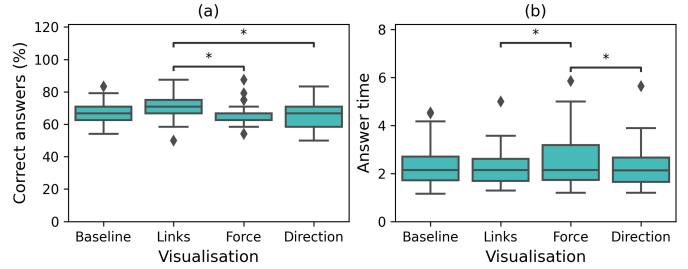


Fig. 3. Diagrams showing medians and quartiles of (a) Correct answer rate in the first task according to the visualisation displayed (b) Answer time in the first task according to the visualisation displayed

be outliers using the interquartile range method. Consequently, the following tests were performed on the results obtained from 45 participants.

1) *Impact of visualisations:* When analysing the effect of the visualisations on the rate of correct answers from participants in anticipating fragmentation, the Friedman test shows statistically significant differences ($Q = 13.67, p = 0.003$). The Nemenyi post-hoc test shows that visualisation displaying links was found to be significantly more helpful in anticipating fragmentation than either dominant force ($p = 0.011$) or direction ($p = 0.015$). Figure 3.a shows the obtained 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 visualisation displaying the dominant force was found to significantly increase the answer time in anticipating fragmentation than either visualisation displaying links ($p = 0.028$) or direction ($p = 0.022$). Figure 3.b shows the obtained results.

2) *Impact of swarm behaviour:* When analysing the effect of the swarm's behaviours on the rate of correct answers from participants 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 4.a 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 in anticipating fragmentation than either expansion ($p < 0.001$) or swarming ($p = 0.013$). Figure 4.b shows the obtained results.

B. Results of the appropriate control task

While analysing the results of the second task, 2 participants were removed from the analysis, as they were considered to be outliers using the interquartile range method. Consequently, the following tests were performed on the results obtained from 49 participants. In this analysis, it was verified if the visualisation and the behaviour displayed by the swarm have

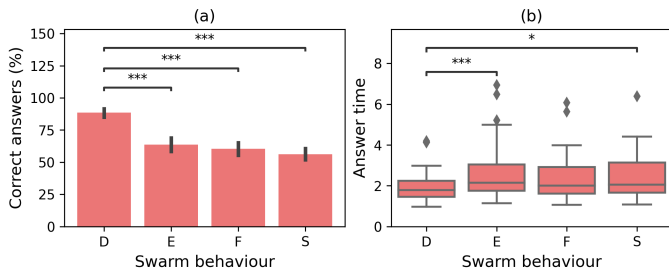


Fig. 4. (a) Correct answers rate in the first task according to the behaviour displayed by the swarm, by showing mean values and standard deviations (b) Answer time in the first task according to the behaviour displayed by the swarm, by showing medians and quartiles (D) Densification (E) Expansion (F) Flocking (S) Swarming

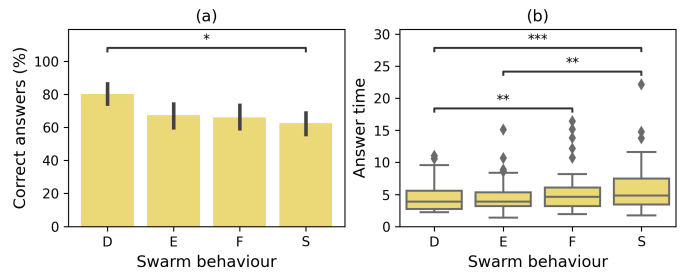


Fig. 6. (a) Correct answers rate in the second task according to the behaviour displayed by the swarm, by showing mean values and standard deviations (b) Answer time in the second task according to the behaviour displayed by the swarm, by showing medians and quartiles (D) Densification (E) Expansion (F) Flocking (S) Swarming

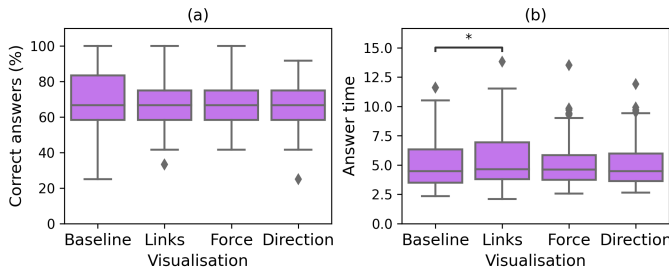


Fig. 5. Diagrams showing medians and quartiles of (a) Correct answers rate in the second task according to the visualisation displayed (b) Answer time in the second task according to the visualisation displayed

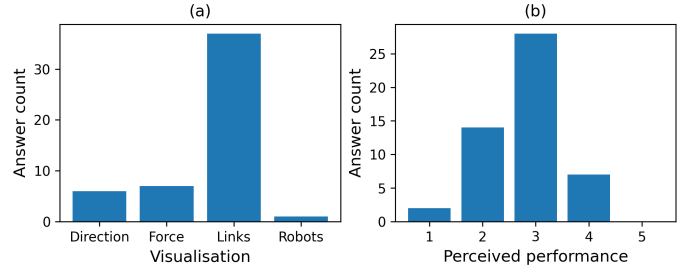


Fig. 7. (a) Visualisation most useful for the first task according to participants (b) Participants’ opinion of their success in the first task, on a scale of 1 (Not at all) to 5 (Perfectly)

an impact on the human capability to choose the appropriate control to prevent fragmentation.

1) *Impact of visualisations:* When analysing the effect of visualisations on the rate of correct answers from participants in choosing the appropriate interaction, the Friedman test shows no statistically significant differences ($Q = 5.80, p = 0.12$). Figure 5.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 visualisation displaying the links was found to significantly increase the answer time in choosing the right interaction than the baseline ($p = 0.015$). Figure 5.b shows the obtained results.

2) *Impact of swarm behaviour:* When analysing the effect of the swarm’s behaviours on the rate of correct answers from participants 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 6.a 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 6.b

shows the obtained results.

C. 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.

1) *Feedback on the first task:* Regarding the distribution of the preferred visualisations for the first task, the visualisation displaying links has the highest proportion (73%), compared to displaying direction (12%), dominant force (14%) or the baseline (2%). Figure 7.a shows the obtained results. Concerning participants’ opinions of their performance in the first task, the mean value is 2.78 (median = 3). Figure 7.b shows the obtained results.

2) *Feedback on the second task:* Regarding the distribution of the preferred visualisations for the second task, the visualisation displaying the dominant force has the highest proportion (47%), compared to displaying links (25%), direction (24%) or the baseline (4%). Figure 8.a shows the obtained results. Concerning participants’ opinions of their performance in the second task, the mean value is 3.03 (median = 3). Figure 8.b shows the obtained results.

3) *Controls understanding:* Two principles were used in the second task controls, cohesion and alignment. Participants

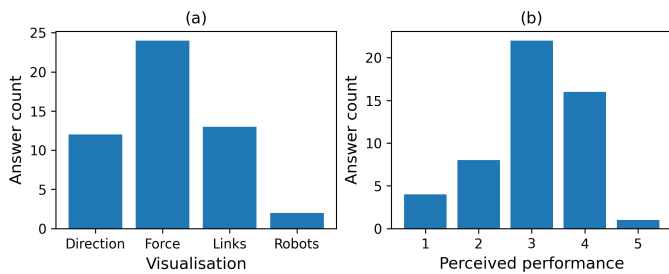


Fig. 8. (a) Visualisation most useful for the second task according to participants (b) Participants' opinion of their success in the second task, on a scale of 1 (Not at all) to 5 (Perfectly)

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.

VI. DISCUSSION

The results presented in section V show that the visualisation displaying the links between robots scored significantly higher on the anticipation task than the other two, direction and dominant force, and was considered the most useful by the participants. For the appropriate control task, there is no significant difference between the three visualisations, and half of the participants rated the visualisation showing the dominant forces as the most useful. In both tasks, none of the proposed visualisations improved the results compared with the baseline showing only robots. Concerning the swarm behaviours, results show that fragmentation resulting from the densification behaviour is easier to anticipate than either expansion, flocking or swarming. For the appropriate control

task, fragmentation resulting from the densification behaviour was easier to correct than swarming.

A. Visualisation of links

The results obtained by the visualisation displaying links between interacting robots are different from those expected. This visualisation was expected to be more helpful to the users than the two other visualisations and the baseline for anticipating fragmentation, as it was designed to show just how close two agents are to breaking up. The quantitative results show that this visualisation enables users to anticipate fragmentation more accurately than the display of directions and forces. Moreover, 73% of the participants said they preferred this visualisation for anticipating fragmentation. However, there is no significant difference between displaying links and displaying nothing on robots in terms of correct answers rate. A comment from a participant could help explain the difference between the qualitative and quantitative results. This comment says 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* [39] to identify robots that are part of the same group. Looking at the causes of the participants' wrong answers for each visualisation to check whether this effect is visible in the results, no significant difference appears. This suggests that this visualisation only conveys the connectivity of the swarm, but not the dynamics of their interaction. A visualisation focusing on the dynamics of the interaction between the robots might be more appropriate in this context, perhaps using calculated information rather than displaying raw information. Concerning the second task, this visualisation does not therefore seem to be useful for choosing the appropriate control and preventing fragmentation. But using the same reasoning as for the first task, a visualisation focusing on the dynamics of the interaction between the robots could convey useful information to users.

B. Visualisation of direction

The quantitative results obtained by the visualisation displaying the directions of the robots show no difference with the display of the robots alone for both tasks. In general, displaying directions is useful for representing movement when it is absent or imperceptible [45]. However, given that humans are naturally capable of perceiving the movement of moving elements, displaying the direction was not expected to provide the users with any more information than they already perceive. Consequently, in the context of fragmentation, the results suggest that displaying the direction of the robots does not improve humans' ability to anticipate fragmentation, nor their ability to choose the right interaction from among several proposals. However, although this visualisation appears to display information already perceived by humans [43], transmitting other information about the action aspect of the swarm other than direction (e.g. the relative movement of the robots compared to the swarm) could give users the

information they need to understand the dynamics of the swarm.

C. Visualisation of dominant force

The visualisation showing the dominant force of the robots aims to help users understand the swarm model by displaying information related to the model rather than explaining its functioning. This visualisation displays information about the dominant force of the robots, and as stated by [8], it might aid comprehension of the model. As the second task aims to modify the intensity of these forces, it could be expected that it also helps to choose the right interaction. However, the quantitative results obtained by the visualisation displaying the dominant force show no difference with the display of the robots alone for both tasks. Moreover, 5 participants commented that they were disturbed by this visualisation. Systematically displaying the dominant force of robots turns out to be counter-productive, as it becomes illegible when the swarm exhibits certain dynamics, such as swarming. The displayed arrows are pointing in all directions, are coloured differently within the same neighbourhood, and all change at the next time step, making it difficult to select the area to focus on and analyse the swarm in a reasonable amount of time. Information acquisition must be immediate, especially when fragmentation can occur in a matter of seconds, and on swarms of much more than 40 robots. This complicates the design of a visualisation based on the swarm model. The visualisation must be detailed enough to represent the swarm model, but simple enough for users to understand [46]. Taking the example of the display of dominant forces, it could be simplified by displaying the information conditionally, by choosing only certain arrows to be displayed, but also by dynamically aggregating some of the information to reduce the amount of information transmitted [8]. Otherwise, rather than displaying the dominant forces, other visualisations representing the decision aspect of the robots could be used, more or less transformed.

D. Swarm Behaviours

Among the four behaviours compared in this study, densification is significantly easier to anticipate. It obtained 88% of correct answers on average, while the others obtained an average of around 60% of correct answers. In addition, the densification behaviour is also the one with the best score for choosing the right interaction, with an average of 80%. However, there was no significant difference with the other categories, which scored around 65%. A difference is noticeable in comparing these results with those of [6]. In their study, participants were free to anticipate or not their response to identify fragmentation. In the end, they 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 [6], as fragmentation in densification behaviour proved easier to identify. The authors suggest that the notion of *Common Fate* from *Gestalt theory* [39] 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 [6], and the use of an immersive environment, the results are consistent.

E. Perception Decision Action loop

The results obtained in the first task suggest that it is preferable to focus the choice of visualisation on the perception aspect of the robots' PDA loop rather than on the action or decision aspect to help users anticipate fragmentation. The visualisation of the links did not improve the results compared with the baseline, probably because it did not sufficiently convey the dynamics of the interactions, which would have made it possible to anticipate future breaks in the swarm, as detailed in section VI-A. In contrast, displaying the direction to represent the action aspect of the swarm conveys information that is already perceived, which has no impact on participants' choices, as the results show. Instead of directions to represent the action aspect, visualisation should focus on information that is not perceived by humans, by improving contrast or transforming low-level information into high-level information, for example. Finally, the systematic display of the dominant forces of the robots to represent the decision aspect was probably too complex for the participants. The design of such a visualisation must consider the users' expertise, the human capability of perception, the task to be accomplished and the complexity of the model to extract only the necessary information. Another interpretation of the results on dominant forces lies in the notion of mental model, which consists of an internal and simplified representation of a dynamical system, allowing humans to describe, explain and predict its form, function and state [47]. According to the qualitative results, it seems that a majority of participants understood the effects of the model's core principles on collective behaviour and are able, to some extent, to describe and explain it. However, the design of the second task requires another *level* of comprehension, as participants have to also be able to predict controls' effects on the robots. Future works will need to investigate the influence of the visualisation of the dominant force, as it may have contributed towards the understanding of

how the model works, but not enough to allow the prediction of how the model will react.

F. Limitation and future work

Gathering enough expert swarm operators is very difficult, so we had to experiment on people with no prior knowledge of swarms. The results could therefore be different with experienced users. Moreover, the results of the second task showed that the proposed visualisations did not affect participants' choice of appropriate control, as no difference is noticeable. As the controls were deliberately imposed on the participants, limited in choice, sometimes counter-intuitive and no feedback on their effect was shown to the participants, it is possible that the task did not allow any differences to be highlighted.

By observing the differences in results within the same category of videos, we can note various parameters that could influence the users' answers during the anticipation task: the speed of expansion of the swarm, its final expansion, the presence of empty space and the presence of centres of gravity. It would then be interesting to evaluate the impact of these parameters. In addition, this study shows elements that could influence the design of future visualisations applied to Human-Swarm Interaction. Future work should consider conducting additional studies to determine what information in the swarm is perceived or not by users, enabling the design of visualisations that do not convey information already perceived, but instead transmit information users are lacking. Finally, future work should also consider exploring new visualisations for conveying the necessary information, taking into account the effect that can be induced by the number of robots making up the swarm, notably related to the *Gestalt theory* principles, which can increase cognitive overload or be counter-productive.

VII. CONCLUSION

In this document, the impact of three visualisations on the ability of humans to prevent fragmentation is evaluated, by presenting 51 participants with two distinct tasks. These respectively assessed the ability to anticipate and the ability to choose the right interaction to avoid fragmentation, using videos containing swarms that fragment or not. The localised and systematic display of robot interactions using links, robot direction using arrows and dominant robot strength using coloured arrows does not help anticipate or choose the right interaction to prevent fragmentation from occurring. Finally, humans more easily anticipate and prevent fragmentation in densification behaviour than in expansion, flocking and swarming. These results show that there is a real need to help humans prevent fragmentation, notably for flocking, expansion and swarming behaviours. Furthermore, the results show that the proposed visualisations do not improve the ability of the users to prevent fragmentation and that it is therefore necessary to explore new solutions. We suggest exploring solutions conveying information that is not already transmitted by the swarm, that does not make swarm monitoring more complicated than it already is, perhaps using conditional or global visualisation.

REFERENCES

- [1] E. Şahin, "Swarm robotics: From sources of inspiration to domains of application," in *Swarm Robotics*, E. Şahin and W. M. Spears, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 10–20.
- [2] C. W. Reynolds, "Flocks, herds and schools: A distributed behavioral model," in *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '87. New York, NY, USA: Association for Computing Machinery, 1987, p. 25–34. [Online]. Available: <https://doi.org/10.1145/37401.37406>
- [3] S. Chen, D. Yin, and Y. Niu, "A survey of robot swarms' relative localization method," *Sensors*, vol. 22, no. 12, p. 4424, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/12/4424>
- [4] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: algorithms and theory," *IEEE Transactions on Automatic Control*, vol. 51, no. 3, pp. 401–420, 2006.
- [5] J. Scholtz, "Theory and evaluation of human robot interactions," in *36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the*. IEEE, 2003, pp. 10–pp.
- [6] 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*, ser. HRI '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 250–258. [Online]. Available: <https://doi.org/10.1145/3610977.3634961>
- [7] 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)*, 2022, pp. 11 258–11 264.
- [8] A. Kolling, P. Walker, N. Chakraborty, K. Sycara, and M. Lewis, "Human interaction with robot swarms: A survey," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 9–26, 2016.
- [9] M. Wooldridge and N. R. Jennings, "Intelligent agents: theory and practice," *The Knowledge Engineering Review*, vol. 10, no. 2, p. 115–152, 1995.
- [10] 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)*. IEEE, 2016, pp. 002 436–002 441.
- [11] J. Harvey, K. E. Merrick, and H. A. Abbass, "Assessing human judgment of computationally generated swarming behavior," *Frontiers in Robotics and AI*, vol. 5, p. 13, 2018. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/frobt.2018.00013>
- [12] G. Punzo, J. Simo, D. J. Bennet, and M. Macdonald, "Characteristics of swarms on the edge of fragmentation," *Phys. Rev. E*, vol. 89, p. 032903, Mar 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevE.89.032903>
- [13] L. G. Jeston-Fenton, S. Abpekar, and K. Kasmarik, "Visualisation of swarm metrics on a handheld device for human-swarm interaction," in *2022 26th International Conference Information Visualisation (IV)*, 2022, pp. 149–154.
- [14] 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*, vol. 23, no. 5, pp. 562–587, 2022. [Online]. Available: <https://doi.org/10.1080/1463922X.2021.1983887>
- [15] A. Capiola, D. Johnson, I. a. Hamdan, J. B. Lyons, and E. L. Fox, "Detecting swarm degradation: Measuring human and machine performance," in *Virtual, Augmented and Mixed Reality*, J. Y. C. Chen and G. Fragomeni, Eds. Cham: Springer Nature Switzerland, 2023, pp. 325–343.
- [16] 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*, ser. CHI '22. New York, NY, USA: Association for Computing Machinery, 2022. [Online]. Available: <https://doi.org/10.1145/3491102.3517719>
- [17] 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.*, vol. 12, no. 4, jul 2023. [Online]. Available: <https://doi.org/10.1145/3597623>
- [18] 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.

- [19] 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*, 2014, pp. 1195–1201.
- [20] 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)*, 2016, pp. 002 521–002 526.
- [21] 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*, vol. 62, no. 1, pp. 287–291, 2018. [Online]. Available: <https://doi.org/10.1177/1541931218621066>
- [22] 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*, ser. ACE '11. New York, NY, USA: Association for Computing Machinery, 2011. [Online]. Available: <https://doi.org/10.1145/2071423.2071436>
- [23] A. M. R. Denniss, T. E. Gorochowski, 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)*, 2019, pp. 1–6.
- [24] K. Libby, L. Zhong, and N. Van Stralen, "Diagnosing robotic swarms 2 (dr. swarm2)," *Retrieved December*, vol. 13, p. 2020, 2020.
- [25] 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*, vol. 5, 2018. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/frobt.2018.00087>
- [26] M. Chen, P. Zhang, Z. Wu, and X. Chen, "A multichannel human-swarm robot interaction system in augmented reality," *Virtual Reality & Intelligent Hardware*, vol. 2, no. 6, pp. 518–533, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2096579620300905>
- [27] S. Batra, J. Klingner, and N. Correll, "Augmented reality for human-swarm interaction in a swarm-robotic chemistry simulation," *Artificial Life and Robotics*, vol. 27, no. 2, pp. 407–415, 2022.
- [28] P. M. Walker, "Improving operator recognition and prediction of emergent swarm behaviors," Ph.D. dissertation, University of Pittsburgh, 2017.
- [29] 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*, ser. ASONAM '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 1275–1279. [Online]. Available: <https://doi.org/10.1145/2492517.2500291>
- [30] D. Miner and N. Kasch, "Swarmvis: a tool for visualizing swarm systems," *UMBC Computer Science*, vol. 636, 2008.
- [31] 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)*, 2019, pp. 1718–1723.
- [32] 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)*, 2012, pp. 3009–3014.
- [33] 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*, ser. CHI EA '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 3443–3448. [Online]. Available: <https://doi.org/10.1145/1520340.1520500>
- [34] 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*, vol. 11, no. 4, pp. 302–311, 2018. [Online]. Available: <https://doi.org/10.9746/jcmsi.11.302>
- [35] —, "Navigatorch: Projection-based robot control interface using high-speed handheld projector," in *SIGGRAPH Asia 2019 Emerging Technologies*, ser. SA '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 31–33. [Online]. Available: <https://doi.org/10.1145/3355049.3360538>
- [36] 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*, 2003, pp. 6 pp.–.
- [37] 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)*, 2018, pp. 422–427.
- [38] 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)*, 2019, pp. 7355–7361.
- [39] 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*, vol. 138, no. 6, p. 1172, 2012.
- [40] J. D. Hasbach, T. E. F. Witte, and M. Bennewitz, "On the importance of adaptive operator training in human-swarm interaction," in *Adaptive Instructional Systems*, R. A. Sottolare and J. Schwarz, Eds. Cham: Springer International Publishing, 2020, pp. 311–329.
- [41] E. R. Hunt, S. Jones, and S. Hauert, "Testing the limits of olfactory stigmergy in high-density robot swarms," *Royal Society open science*, vol. 6, no. 11, p. 190225, 2019. [Online]. Available: <https://doi.org/10.1098/rsos.190225>
- [42] 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*, ser. ASONAM '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 1275–1279. [Online]. Available: <https://doi.org/10.1145/2492517.2500291>
- [43] S. N. Watamaniuk and A. Duchon, "The human visual system averages speed information," *Vision research*, vol. 32, no. 5, pp. 931–941, 1992.
- [44] 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*, vol. 94, pp. 761–775, 2019. [Online]. Available: <https://doi.org/10.1007/s10846-018-0866-9>
- [45] 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*, vol. 65, pp. 311–320, 1991.
- [46] C. L. Remondino, B. Stabellini, P. Tamborini *et al.*, "Visualizing complex systems," in *Visualizing Complex Systems*. Politecnico di Torino, 2019, pp. 17–31.
- [47] N. A. Jones, H. Ross, T. Lynam, P. Perez, and A. Leitch, "Mental models: an interdisciplinary synthesis of theory and methods," *Ecology and society*, vol. 16, no. 1, 2011.