



Toward the Legibility of Multi-robot Systems

BEATRICE CAPELLI, Department of Sciences and Methods for Engineering, University of Modena and Reggio Emilia, Reggio Emilia, Italy and System Logistics, Fiorano Modenese, Italy

MARÍA SANTOS, Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, USA

LORENZO SABATTINI, Department of Sciences and Methods for Engineering, University of Modena and Reggio Emilia, Reggio Emilia, Italy

Communication is crucial for human-robot collaborative tasks. In this context, legibility studies movement as the means of implicit communication between robotic systems and a human observer. This concept has been explored mostly for manipulators and humanoid robots. In contrast, little information is available in the literature about legibility of multi-robot systems or swarms, where simplicity and non-anthropomorphism of robots, along with the complexity of their interactions and aggregated behavior impose different challenges that are not encountered in single-robot scenarios. This article investigates legibility of multi-robot systems. Hence, we extend the definition of legibility, incorporating information about high-level goals in terms of the coordination objective of the group of robots, to previous results that focused solely on the legibility of spatial goals. A set of standard multi-robot algorithms corresponding to different coordination objectives are implemented and their legibility is evaluated in a user study, where participants observe the behavior of the multi-robot system in a virtual reality setup and are asked to identify the system's spatial goal and coordination objective. The results of the study confirmed that coordination objectives are discernible by the users, hence multi-robot systems can be controlled to be legible, in terms of spatial goal and coordination objective.

CCS Concepts: • **Human-centered computing** → **Interaction paradigms; Interaction design theory, concepts and paradigms**

Additional Key Words and Phrases: Multi-robot systems, human-centered robotics, swarms, human-aware multi-robot systems

ACM Reference Format:

Beatrice Capelli, María Santos, and Lorenzo Sabattini. 2024. Toward the Legibility of Multi-robot Systems. *ACM Trans. Hum.-Robot Interact.* 13, 2, Article 21 (June 2024), 32 pages. <https://doi.org/10.1145/3647984>

1 INTRODUCTION

As robotic platforms become increasingly affordable, reliable, and efficient, multi-robot systems are expected to slowly infiltrate many activities in our everyday lives, blurring the divide between human and robot spaces. The inherent redundancy of multi-robot systems and robotic

Authors' address: B. Capelli, Department of Sciences and Method for Engineering, University of Modena and Reggio Emilia, via Amendola 2, 42122 Reggio Emilia, Italy and System Logistics, Via Ghiarola Vecchia 50, 41042 Fiorano Modenese, MO, Italy; e-mail: beatrice.capelli@unimore.it; M. Santos, Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08540, USA; e-mail: maria.santos@princeton.edu; L. Sabattini, Department of Sciences and Methods for Engineering, University of Modena and Reggio Emilia, Reggio Emilia, Italy; e-mail: lorenzo.sabattini@unimore.it



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2024 Copyright held by the owner/author(s).

ACM 2573-9522/2024/06-ART21

<https://doi.org/10.1145/3647984>

swarms, afforded by their multiple, and often interchangeable, individual components and their spatial flexibility will likely increase their deployment in a variety of applications, including precision agriculture, logistics, environmental monitoring, or surveillance, among others (see, e.g., Reference [60] for a survey).

While an efficient and safe coordination of the robots is imperative for the fulfillment of each specific task objective, the intersection of robot and human work-spaces demands effective communication among them. The need for such communication, which has become a prominent subject for the interaction between humans and individual robots (see Reference [21] for an example with manipulators), poses further challenges for the multi-robot scenario due to the complexity associated with the coordination among numerous individual robots. This hurdle becomes apparent in both communication directions: how can a human user exert control over the multi-robot system and how can the robots give some information to the human? While some authors have addressed the former question (e.g., References [28, 51]), literature addressing the latter is lacking.

Intention communication is the key to a fruitful interaction between robots and humans [27]. Therefore, in this article, we focus on *legibility* of multi-robot systems. In particular, we consider legibility as the ability of a robot, or in this case of a group of robots, to manifest its intention without the use of explicit communication. To this end, this work studies a series of standard coordinated control strategies and their adequacy in conveying spatial goals of the multi-robot system to a human observer. Furthermore, the different algorithms are analyzed to determine if a human observer is capable of distinguishing the coordination objective of the multi-robot system, defined by each specific decentralized coordinated control algorithm, and how different parameters that characterize the control algorithm may affect the interpretability of the multi-robot team motions.

The article is organized as follows. The remainder of Section 1 provides a literature review on different aspects related to legibility of robotic systems and outlines the main contributions of this article. The definition of legibility for multi-robot systems is presented in Section 2. Section 3 proceeds to describe the coordination objectives whose legibility is investigated in the article. Such investigation is conducted via a user study, whose procedure is outlined in detail in Section 4. Then, in Section 5, we report the results of the user study and their analysis. Main limitations of the achieved results are discussed in Section 6, while discussion and conclusions are provided in Section 7.

1.1 Literature Review

1.1.1 Proximal Interaction. When considering **Human-Robot Interaction (HRI)**, a key factor is the differentiation between proximal and remote interaction [36], i.e., the users share the same environment with the robot or they do not, respectively. In applications with remote interaction, the user and the robotic system are in different physical locations [52], e.g., search and rescue operations [50], aerospace repairs [30], bomb disposals [59], or surgical procedures [47]. This type of interaction is characterized by a user-robot communication by means of a physical instrument, such as a teleoperation tool or a joystick, that allows the user to provide the robot with input commands and, at the same time, allows the user to receive information from the robot, in terms of visual, audio, or force feedback.

Instead, in proximal interaction, the environment is shared by the human user and the robotic platform, and the communication is carried out directly through, e.g., gesture recognition, speech, and gaze engagement. This condition of proximal interaction will be the focus of this article. As the robots share the same environment with the user, in addition to communication, we need to take into account the safety of the user. For example, in physical-HRI, the robot, usually a manipulator, has to safely interact with the user, which is achieved by imposing velocity limits as a function of

the distance to the user [25]. Even though safety is an issue in proximal interaction applications, in this work, we will not address this aspect due to the fact that the user study was carried out in virtual reality, and that the focus was mainly on legibility.

1.1.2 Implicit Communication. From the point of view of the communication, the bi-lateral question still arises: how should the user communicate to the robot and how should the robot itself refer important information back to the user? This work focuses on the latter question. Along the lines of Reference [39], the approaches can be divided into explicit and implicit cues, based on the use of ad hoc apparatus for communication purposes. For example, explicit cues that have been used in the literature include LEDs [66], light indicators [45], voice [62], gaze movement of a humanoid robot [45], or **Augmented Reality (AR)** [12, 29, 56]. Although similar to a user interface, AR has the advantage of merging the information directly into the real sight of the user. However, the presence of a tool (e.g., smart glasses or a tablet), places an additional entity in the communication channel: there is no longer direct communication between the robot and the user. Furthermore, using an additional tool may not always be possible, due to task-related constraints (e.g., a tablet cannot be used, because the user is wearing gloves or is using some tool) or may require some specific training (e.g., the user needs to learn how to use smart glasses). The need for task-specific training also emerges with the use of explicit cues linked to colors or shapes, as some sort of alphabet is needed to match the cues to their meanings. The aforementioned considerations become even more impactful in the case of multi-robot systems: the cognitive complexity, defined in Reference [36] as the amount of cognitive resources necessary for a user to understand and provide input to a multi-robot system, significantly increases with the number of robots. Furthermore, considering the simple mobile robots typically adopted in multi-robot applications, the lack of anthropomorphism impedes the use of cues mimicking human communication. For instance, using voice cues would imply choosing which robot communicates, what information needs to be delivered, and how the communication channel is established.

Implicit cues, conversely, use directly the movement of the robots to convey information, thus leveraging human understanding of movement and the tendency of people to discover intentionality in everything, even in robots [16]. Along these lines, the authors of Reference [65] proposed the use of animation principles to communicate intent in assistive free-flyers. By manipulating the velocities and trajectories of the UAVs, they discovered how to effectively improve communication and that the users preferred motions with ad-hoc modification to infer some significance. The work presented in References [19–21] focused on manipulators and studied their legibility, i.e., their ability to communicate intent through movement, to improve collaborative tasks. In Reference [21] the authors investigated the difference between legibility and predictability, in Reference [19] they studied how to generate legible motions, and in Reference [20] they verified the utility of these legible motions into a collaborative task. In their studies the motion of the robot is generated by means of a planner that optimizes legibility itself. We took inspiration from their definition of legibility to proceed in our study toward multi-robot applications. Hence, preliminary works were presented in References [5, 9, 10], where the characteristics of the motion of a multi-robot system (see Section 1.1.3 for more details) were directly studied. However, no planner to generate legible motions was defined. Furthermore, high-level coordination objectives were not considered in these previous works, and are explicitly addressed in this article. Such high-level objectives may be linked to the studies on the communication of robots' plan, which is a topic studied in the **Artificial Intelligence (AI)** field [12, 38, 75]. In this field, in particular in the context of **Human-Machine Interface (HMI)** and with mechanical interfaces, such as in teleoperation tasks, legibility is often referred to as *transparency*. In Reference [11] a detailed description of the differences and of the implementations of legibility and transparency in AI can be retrieved.

1.1.3 Legibility in Multi-robot Systems. In the context of social applications, robots must effectively interact with people in an engaging manner [4], a requirement that oftentimes demands expressing emotion through their actions. In robotic swarms and multi-robot systems composed by simple, non-anthropomorphic robots, conveying emotions to an observer can pose an increased difficulty. In these cases, the emotional expressivity and readability of the movement of the robotic swarm has been linked to different parameters related to the individual control of the robots in the swarm and the inter-robot interactions [18, 42], under the context of standard collective behaviors such as cyclic pursuit or coverage [58] or relating them to dancing [63]. Regularity of the motion pattern and different levels of cohesion in the group of robots have been identified as important aspects in human understanding, as discussed in References [32, 61]. Building upon these ideas, in this article, we investigate the influence of the choice of control parameters on the capability, for the user, to infer spatial and coordination objectives of the multi-robot systems.

Legibility of a multi-robot system was introduced in our previous works [9, 10] as the ability of the multi-robot system to communicate to the user its spatial goal. Such ability was characterized based on two performance indices: (1) the correctness of the communication and (2) the time needed for this communication. The first question to answer was whether the motion of a multi-robot system was able to communicate some information to the user. We investigated this issue in our previous work [10], where we found out that some *motion-variables* (namely, the trajectory, the dispersion and the stiffness of the multi-robot system) were statistically relevant for the communication of the spatial goal of the group. The *motion-variables* were chosen to describe some important characteristics of the multi-robot system. The user study was conducted in a virtual reality setup, where the user was able to share the same environment with the robots. The kind of controller used to define the movement of the robots was based on artificial potential fields.

The following step was to understand if the same *motion-variables* could be used to communicate the same information, i.e., the spatial goal, of multiple groups of robots at the same time [9]. The results of this study were similar to those of the previous user study and the *motion-variables* confirmed to be statistically relevant with respect to the legibility of a multi-robot system.

In Reference [34] the legibility of multi-robot systems was further investigated, together with glanceability, defined as the ability to communicate a spatial goal throughout the whole movement of the robots. The authors focused on different types of control laws, similar to those implemented in this study, but they investigated also the effect of target difficulty, namely, whether the spatial goals had the same number of neighboring goals or not, and initial radius, namely, how much the robots were spread out at the beginning of the experiment. Only the legibility with respect to the spatial goal was analyzed, with rendezvous (i.e., consensus) being the most legible controller and the trajectory-based motion the best for glanceability. Along similar lines, a user study was performed in Reference [69] to understand if humans are able to identify aggregation/dispersion behaviors in swarms of robots. In both cases, the user studies were carried out on a standard screen and with a top-down view that is not suitable with our idea of proximal interaction. In fact, the user study reported in this article, and in the previous works [9, 10], were conducted in a virtual reality setup with the objective of imitating as closely as possible a proximal interaction scenario.

1.2 Statement of Contributions

This article studies the legibility of multi-robot systems. In particular, the contributions of this article are:

- Extended definition of legibility for multi-robot systems;
- Extended characterization of the legibility of spatial goals considering standard control laws through a user study;

- Characterization of the legibility of coordination objectives, defined in terms of standard decentralized coordinated control strategies, in the user study.

Taking into account the results of previous works [9, 10], we extend the study on the legibility with respect to spatial goals. In particular, we consider the effect of different coordination objectives and of different numbers of robots on the legibility of the spatial goal. To this end, we identify a set of standard decentralized coordinated control laws, used extensively in the multi-robot systems literature, for realizing desired coordination objectives. We then investigate the legibility characteristics of each of these control laws. To the best of the authors' knowledge, this is the first work to investigate the legibility of coordination objectives.

Our previous advances on the legibility of multi-robot systems [5, 9, 10] focused on a microscopic approach, adding to the trajectory parameters typically analyzed in legibility studies [21] those parameters affecting inter-robot relationships, as in Reference [41]. On the contrary, this article adopts a holistic view of multi-robot legibility, thus relating to the observation of the aggregate movement of the team.

2 DEFINITION OF LEGIBILITY IN MULTI-ROBOT SYSTEMS

What does legibility mean for a multi-robot system? Can we define it in terms of *spatial goal* (where are the robots going) and *coordination objective* (what are the robots doing)? These can be useful when sharing the space with the user. Legible motion endows robots with the ability to communicate information directly, without the need of a specific user interface. While user interfaces continue to be the standard in many human-robot interaction scenarios, here we want to understand if it is possible to communicate implicitly this information. The prototype coordinated control algorithms that we choose to investigate for the definition of the coordination objectives are: consensus, formation, and coverage. We chose these controllers due to their ubiquitous use in multi-robot systems: as pointed out in a recent survey paper [13], these control strategies represent the fundamental building blocks that are generally exploited in the literature to address most of the multi-robot control problems.

Given that the legibility of multi-robot systems is still in its early stages, we need to understand what legibility means in this context. We consider legibility as the ability of a group of robots to communicate to the user its spatial goal and/or its coordination objective. Hence, we propose the following definition:

Definition 1. [Legible multi-robot system] A multi-robot is legible if its spatial goal and coordination objective are understandable by a user by solely observing the motion pattern (i.e., without the use of explicit communication methods), where

- The *spatial goal* represents the overall direction of motion of the multi-robot system;
- The *coordination objective* represents how robots are moving with respect to each other and how they are aggregating.

3 MULTI-ROBOT COORDINATION OBJECTIVES

A large body of work on multi-robot systems and robotic swarms deals with how to coordinate large teams of robots to produce collective behaviors that go beyond the capabilities of the robots as individuals [3]. As such, the aim is to perform collective, emerging behaviors that arise from rules executed by individual robots acting on local information. The fulfillment of objectives is thus typically analyzed from the team-level behavior, rather than from the individual actions of the robots. For instance, robot team performance can be evaluated in terms of overall geometrical deployment of the robots, as discussed in Reference [68].

Taking the taxonomy outlined in Reference [13] as a reference, we analyze the legibility of standard coordinated control algorithms commonly used for achieving multi-robot coordination objectives. Specifically, we examine consensus, formation, and coverage algorithms. These algorithms are posed such that the group's coordination objective is encoded through a cost that represents the performance of the team as a whole, and the individuals cooperate toward that objective by following a direction of the gradient of the cost.

In this section, we outline the control laws executed for each of the coordination objectives utilized in the user study. To this end, let us consider a team of N planar robots,¹ with the position of robot i being denoted as $x_i \in \mathbb{R}^2$, $i \in \mathcal{N} = \{1, \dots, N\}$. While many kinematic configurations can be considered for the individual platforms in robotic swarms (e.g., omnidirectional [35] or differential drive [48, 71, 72]), we here opt for omnidirectional robots to minimally distort the ideal collective behavior considered in the control laws. Consequently, the movement of the robots can be described through single integrator dynamics, i.e.,

$$\dot{x} = u,$$

with $x = [x_1^T, \dots, x_N^T]^T \in \mathbb{R}^{2N}$ being the aggregate state of the system, and $u = [u_1^T, \dots, u_N^T]^T \in \mathbb{R}^{2N}$, its control. If a coordination objective is to be fulfilled, then the control input of each particular robot, u_i , $i \in \mathcal{N}$, will be influenced by the state and the actions of its neighbors. This flow of information among adjacent robots is typically represented through a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where the vertex set, $\mathcal{V} = \{1, \dots, N\}$, represents the robots and the edge set, $\mathcal{E} \in \mathcal{V} \times \mathcal{V}$, encodes the adjacency relationships, i.e., $(i, j) \in \mathcal{E}$ if robot i communicates with or can sense/be sensed by robot j . For the control laws considered in this article, the relationships between robots are bidirectional, namely, $(i, j) \in \mathcal{E}$ if $(j, i) \in \mathcal{E}$, that is, \mathcal{G} is undirected. In addition, they are subjected to change over time: the edge set \mathcal{E} need not be static. For notational convenience, the neighboring set of robot i is denoted as $\mathcal{N}_i = \{j \in \mathcal{N} \mid (i, j) \in \mathcal{E}\}$.

The remainder of this section provides an outline of the different multi-robot coordination objectives analyzed in the user study in Section 4, i.e., consensus, formation and coverage. For each of them, we include a general description of the expected behavior as well as the laws to be executed by the robots. Additional details about implementation, including parameter choices, are reported in Appendix A for the purpose of reproducibility of results.

3.1 Consensus

The consensus protocol deals with the spatial aggregation of agents in a common point through local coordination of neighbors [1, 14, 43]. Assuming the graph describing inter-robot sensing, \mathcal{G} , is connected², the team converges to a common location if each individual executes the control law

$$\dot{x}_i = \kappa \sum_{j \in \mathcal{N}_i} (x_j - x_i), \quad \forall i \in \mathcal{N}, \quad (1)$$

where $\kappa \in \mathbb{R}_+$ is a proportional gain.³ Executing the control law in Equation (1) implies that robot i will be moving in the direction that averages the relative positions of its neighbors, $x_j - x_i$, $\forall j$ adjacent to robot i . Asymptotically, this will make them converge to a point $\bar{x} = x_1(\infty) = \dots =$

¹Note that, while the control laws outlined in this section are generalizable to an n -dimensional space, i.e., $x_i \in \mathbb{R}^n$, $i \in \{1, \dots, N\}$; we present them here with $x_i \in \mathbb{R}^2$ to relate to the planar robots considered in this study.

²A graph \mathcal{G} is connected if, for every pair of vertices in $\mathcal{V}(\mathcal{G})$, there is a path that has them as its end vertices.

³With \mathbb{R}_+ , we denote the set of strictly positive real numbers.




Fig. 1. Evolution of the consensus protocol in Equation (1) on a team of 15 robots. The red circles represent idealized robots with no physical footprint, which allows them to overlap. The black lines represent the underlying graph, i.e., the pairwise connections among the robots. Given that the graph is connected, the consensus protocol makes the robots to converge to a point.

$x_N(\infty)$, that is, the robots will reach consensus about where to meet, as long as the graph stays connected.

An instance of the consensus protocol in Equation (1) is shown in Figure 1, where a group of ideal robots, modeled as points, starts scattered over the domain and agree on a location where to meet. Note that, in the experiments included in Section 4, the robots have a physical footprint and, therefore, are unable to finish stacked on top of each other. As opposed to the ideal case presented in Figure 1, when executing the consensus protocol on a team of real robots, they will aggregate as much as possible without incurring inter-robot collisions, thanks to the implemented collision avoidance (see Appendix A.2).

3.2 Formation Control

A standard problem when controlling a multi-robot team is that of moving the individual robots so that they display a particular shape. This is typically achieved through formation control strategies, e.g., References [31, 54, 73], that specify how the individual robots should move to collectively display a shape.

The specification of the geometry to be realized by the team is key to the operation of a formation controller. To this end, we can specify a 2D formation Δ through a set of relative, desired inter-agent distances,

$$\Delta = \{\delta_{ij} \in \mathbb{R} \mid \delta_{ij} > 0, i, j = 1, \dots, N, i \neq j\}.$$

We assume here that Δ is a feasible formation, that is, that the distances are not conflicting and, thus, the geometry can be realized (see Reference [46] for details on feasible formations).

Having specified the formation, we can recover the structure of the consensus protocol in Equation (1) and make each robot move according to a weighted average of the relative positions of its neighbors,

$$\dot{x}_i = \sum_{j \in \mathcal{N}_i} w_{ij}(\|x_i - x_j\|, \delta_{ij})(x_j - x_i), \quad \forall i \in \mathcal{N}, \quad (2)$$

where the weight w_{ij} is positive if the neighboring robot is further than the desired distance, i.e., if $\|x_i - x_j\| > \delta_{ij}$, and negative otherwise. For the experiments in Section 4, the weights are calculated as

$$w_{ij}(\|x_i - x_j\|, \delta_{ij}) = \frac{\|x_i - x_j\| - \delta_{ij}}{\|x_i - x_j\|}, \quad (i, j) \in \mathcal{E}.$$




Fig. 2. Evolution of the formation control protocol in Equation (2) on a team of 10 robots. The red circles represent idealized robots with no physical footprint and the black lines represent the underlying graph. A connection between robots i and j , represented by a graph edge, corresponds to a desired distance, δ_{ij} , that is to be maintained to achieve the desired geometry. The desired formation, composed by two concentric pentagons, is achieved in panel (c).

Note that, to execute the control law in Equation (2), each robot needs to be cognizant of which robots conform its neighborhood, \mathcal{N}_i , their relative positions, $x_j - x_i$, and which distance is to be maintained with respect to each of them, δ_{ij} , as reflected in Δ . Therefore, under this protocol the robots are no longer anonymous: each robot needs to know the identities of other robots to discriminate which ones it needs to coordinate with.

Figure 2 shows the evolution of a team of 10 robots achieving a formation composed by two concentric pentagons, as specified by a formation Δ . The formation control protocol in Equation (2) ensures that the inter-robot distances are maintained but it does not influence the rotation of the formation with respect to the global frame.

3.3 Coverage Control

Coverage control deals with the problem of distributing robots over a domain $D \subset \mathbb{R}^2$ such that the events of interest are optimally monitored by the team [15]. When covering an area, a standard strategy is to partition it so that each robot only takes responsibility over a portion. A typical way of dividing responsibilities is to make robot i be in charge of those points in the domain that are closer to it than to any other robot. This partitioning is known as a Voronoi tessellation [53],

$$V_i(x) = \{q \in D \mid \|q - x_i\| \leq \|q - x_j\|, \quad \forall j \in \mathcal{N}\},$$

where V_i denotes the Voronoi cell assigned to robot i .

In general, this partition is not uniform across the domain in terms of area assigned to each robot. Indeed, the objective of the coverage problem is to monitor closely areas of higher importance by making the robots concentrate around them, which, in turn, results in differently sized Voronoi cells. The relative importance of a point in the domain typically reflects the probability of occurrence of an event or the concentration of a resource and is described through a spatial field, hereinafter referred to as *density function*, ϕ . An example of this density function is depicted in Figure 3, with *yellow* shades corresponding to higher values of ϕ .

When the density function is solely a function of the points in the domain and does not evolve over time, $\phi : D \mapsto \mathbb{R}_+$, the overall coverage performance of the multi-robot team can be quantified by the locational cost defined in References [15, 22],

$$\mathcal{H}(x) = \sum_{i=1}^N \int_{V_i(x)} \|x_i - q\|^2 \phi(q) dq, \quad (3)$$




Fig. 3. Coverage control from Equation (4) executed by a team of 15 robots. The robots are tasked with covering the domain with respect to a density function, whose values are depicted through a heat map where higher values of the density function are represented in yellow and lower ones in green. As the robots move according to the coverage control law, the team concentrates around areas where the density is higher.

with a lower cost corresponding to a better coverage of ϕ . The multi-robot system can minimize this cost following a direction of descent [15], which lends a continuous version of Lloyd’s algorithm [44],

$$\dot{x}_i = \kappa(c_i(x) - x_i), \tag{4}$$

with $\kappa \in \mathbb{R}_+$ a proportional gain, and $c_i(x)$ the center of mass of the Voronoi cell of robot i ,

$$c_i(x) = \frac{\int_{V_i(x)} q\phi(q)dq}{\int_{V_i(x)} \phi(q)dq}. \tag{5}$$

3.3.1 Time-varying Coverage. In some situations, the importance of the points in the domain may evolve over time. As a result, the coverage strategy outlined in Equations (3)–(5) needs to adapt to reflect possible dynamic changes to the density function. Considering a dynamic density, $\phi : (q, t) \in D \times \mathbb{R}_+ \mapsto \phi(q, t) \in \mathbb{R}_+$, the locational cost can be reformulated as

$$\mathcal{H}(x, t) = \sum_{i=1}^N \int_{V_i(x)} \|x_i - q\|^2 \phi(q, t) dq. \tag{6}$$

The introduction of a dynamic density in the locational cost calls for a modification of the coverage controller. As investigated in Reference [40], one can minimize the locational cost in Equation (6) by setting

$$\dot{x} = \left(I - \frac{\partial c}{\partial x} \right)^{-1} \left(\kappa(c(x, t) - x) + \frac{\partial c}{\partial t} \right),$$

where $c = [c_1^T, \dots, c_N^T]^T \in \mathbb{R}^{2N}$ is the vector containing the centers of mass calculated as in Equation (5) but with respect to the time-varying density function, $\kappa \in \mathbb{R}_+$ is a proportional gain, and $I \in \mathbb{R}^{2N \times 2N}$ is the identity matrix.

4 USER STUDY

In this section, we describe the user study that was conducted to investigate the legibility of multi-robot systems. Following the preceding studies in References [9, 10], this user study investigates how the different controllers described in Section 3, as well as how the number of robots executing them, affected the legibility of the system, both in terms of spatial goal and coordination objective. Hence, this study takes a step further in understanding the legibility of multi-robot systems by analyzing how distinguishable the different coordination objectives are for the human observer. To this end, users were asked to recognize the spatial goal of the multi-robot system and the different

kinds of controllers, namely, we wanted to discover if the users can understand the coordination objective of the multi-robot system, in addition to the spatial goal. A representative trial of the experiment is shown in the video available at Reference [8].

4.1 Experimental Scenario

We carried out a user study in a virtual reality setup motivated by the possibility of having a high number of robots sharing their environment with the human user, thus replicating a scenario with proximal interaction. About the use of virtual reality as a means to retrieve data for a user-study, we think that the level of realism carried by modern virtual reality headsets is sufficient to acquire plausible data [67]. In addition, virtual reality allows for arbitrarily large environments and number of robots without the disturbances that could arise in real world experiments (e.g., delays, failures, or noisy communication), therefore mitigating exogenous factors that may influence the perception of the different multi-robot coordination objectives. An alternative choice to this setup would have been to show to the users the movement of the robots with a video [34, 65] or a sequence of photos [55], but we believe that the advantages introduced by virtual reality, such as the perception of proximal interaction, exceed the drawback of using an additional means for the experiments.

In particular, the virtual reality setup was made using Unity, a cross-platform game engine, and it was rendered with the Oculus Rift headset.⁴ The setting was a plain environment with ample space, approximately the size of a tennis court (22 m × 11 m). Within the environment, we placed the spatial goals, which were cubes emitting light of different colors (Figure 4(a)). The simplicity of the environment was purposefully designed to avoid creating distractions for the users, allowing them to focus solely on the movement of the robots. We used 20 omnidirectional robots to make the multi-robot system. The robots, depicted in Figure 6, were taken from the standard Unity assets, and are characterized by a 0.1 m radius.


4.2 Experimental Procedure

The experiment was divided in two parts: one related to the spatial goal of the multi-robot team and one to its coordination objective. The overall procedure is summarized in Figure 7, and described in details hereafter.


4.2.1 First Part. Users were asked to recognize as fast as possible the goal toward which the robot team was moving. The answer was given by means of the Oculus Remote (see Figure 5) combined with the direction of the head of the user. Once the users had identified with a sufficient degree of confidence the probable goal, they had to point toward the chosen goal (in the field of view of the user, a point was shown that corresponded to the direction tracked by the system) and press the central button of the Oculus Remote.

4.2.2 Second Part. Users had to recognize both the spatial goal and the coordination objective of the multi-robot system. The coordination objectives were introduced to the users via three representative examples shown in the video available online at Reference [6]. No information about the kind of controllers or other features of the coordination objectives were disclosed to the users, with the objective of not to insert any bias in the experiments. Furthermore, the examples were significantly different from the movements of the second part of the experiment. In fact, the number of robots did not correspond to neither used in the experiments nor the directions of the movements correspond to one of the goals. The examples consisted of a team of eight robots moving right to left in front of the user. After this introduction, the

⁴The developed virtual reality setup is freely available online (executable file and source code) [7].



(a) Example of the sight of the user in the first part of the experiment. The *colored* cubes are the spatial goals to whom the multi-robot system could move.



(b) Panels with the coordination objectives' name in the second part of the experiment. In this part, the user had to understand the spatial goal and the coordination objective of the multi-robot system.

Fig. 4. Virtual reality environment built in Unity and used during the user study. Virtual reality allows us to replicate with a good degree of fidelity a proximal interaction with a multi-robot system [67].



Fig. 5. Oculus Remote for user's input.

second part of the experiment started and the users were asked to recognize the spatial goal and the coordination objective of the group. They were informed that the amount of time was measured separately, namely, they had to give the answers (spatial goal and coordination objective) as soon as they understood one of the two. Each trial ended when the user gave both answers. The answer for the spatial goal was provided in the same way as in the first part of the experiment, instead the chosen coordination objective was specified through the interface in Figure 4(b). No information about the kind of controllers or other features of the coordination objectives were disclosed to the users, with the objective of not to insert any bias in the experiments.




Fig. 6. Omnidirectional robots used in the experiments, taken from the standard Unity asset.

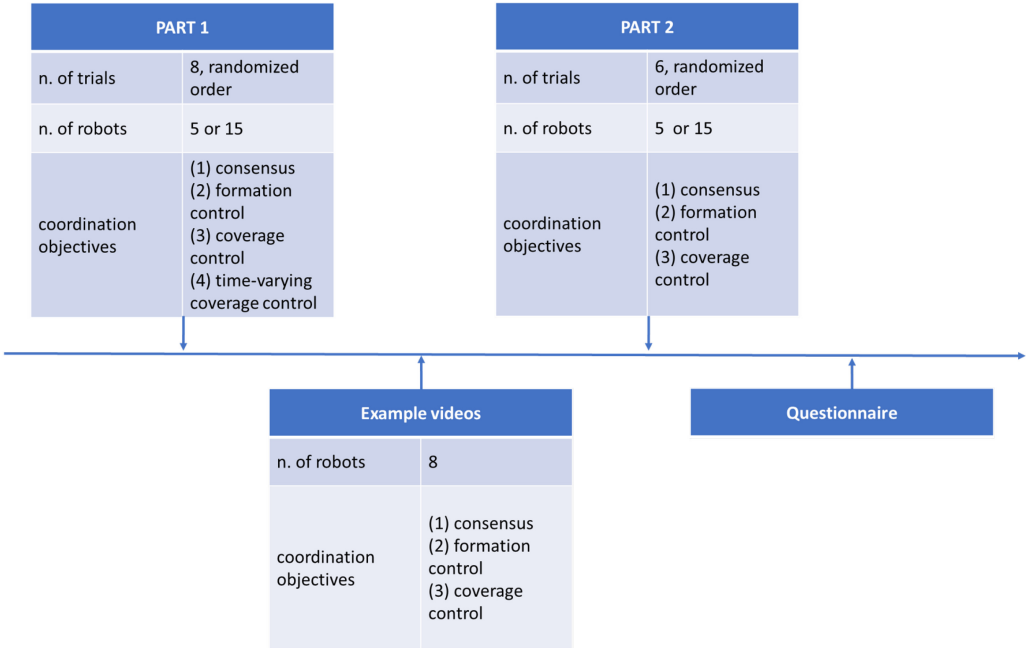


Fig. 7. Summary of the experimental procedure: sequence of activities carried out by the participants.

Furthermore, the examples were significantly different from the movements of the second part of the experiment. In fact, the number of robots did not correspond to neither used in the experiments nor the directions of the movements correspond to one of the goals. The examples consisted of a team of 8 robots moving right to left in front of the user. After this introduction, the second part of the experiment started and the users were asked to recognize the spatial goal and the coordination objective of the group. They were informed that the amount of time was measured separately, namely, they had to give the answers (spatial goal and coordination objective) as soon as they understood one of the two. Each trial ended when the user gave both answers. The answer for the spatial goal was provided in the same way as in the first part of the experiment, instead the chosen coordination objective was specified through the interface in Figure 4(b). No information about the kind of controllers or other features of the coordination objectives were disclosed to the users, with the objective of not to insert any bias in the experiments. Furthermore, the examples were significantly different from the movements of the second part of the experiment. In fact, the number of robots did not correspond to neither used in the experiments nor the directions of the movements correspond to one of the goals. The examples consisted of a team of eight robots moving right to left in front of the user. After this introduction, the second part of the experiment

started and the users were asked to recognize the spatial goal and the coordination objective of the group. They were informed that the amount of time was measured separately, namely, they had to give the answers (spatial goal and coordination objective) as soon as they understood one of the two. Each trial ended when the user gave both answers. The answer for the spatial goal was provided in the same way as in the first part of the experiment, instead the chosen coordination objective was specified through the interface in Figure 4(b). No information about the type of controllers or other features related to the coordination objectives was disclosed to the users, with the objective of avoiding the introduction of any bias in the experiments. Furthermore, the examples were significantly different from the movements of the second part of the experiment. In fact, the number of robots did not correspond to neither used in the experiments nor the directions of the movements correspond to one of the goals. The examples consisted of a team of 8 robots moving right to left in front of the user. After this introduction, the second part of the experiment started and the users were asked to recognize the spatial goal and the coordination objective of the group. They were informed that the amount of time was measured separately, namely, they had to give the answers (spatial goal and coordination objective) as soon as they understood one of the two. Each trial ended when the user gave both answers. The answer for the spatial goal was provided in the same way as in the first part of the experiment, instead the chosen coordination objective was specified through the interface in Figure 4(b).

The experiment ended with a survey with general questions for the user, including demographic information and previous experiences with virtual reality, and some specific questions regarding the experiment itself. While these data do not influence the statistical analysis of the experiment, these questions clarify the perception of the experiments from the users. The questionnaire was based on a 5-point Likert scale.

4.3 Experimental Variables

The *independent variables* of this study are the coordination objectives of the multi-robot system and the number of robots in the group. In the first part, the coordination objectives were: consensus, formation, and coverage. In the case of coverage, both static and time-varying density functions were evaluated with the objective of analyzing whether the two corresponding controllers affect or not the legibility of the multi-robot system. However, in the second part, time-varying coverage was not utilized given that, visually, the robots behaved too similarly to the consensus protocol for this scenario, as seen in preliminary tests. In fact, in both cases the robots agglomerate at the beginning and then they move toward the goal. The only difference was in terms of dispersion in the environment: the time-varying coverage had a higher dispersion with respect to the consensus. Representative examples of trajectories utilized in the experiments are reported in the Appendix A.5.

For the coordination objectives in Section 3, each control law has certain parameters that define the behavior of the multi-robot system. These parameters are distinctive of the controllers and they should not be considered as independent variables. The choice of the parameters is detailed in Appendix A.4. Note that, differently from the previous works [9, 10], *trajectory*, which is the path followed by the group of robots during the movement, is not an independent variable in this study. In fact, based on the results of References [9, 10], the best trajectory to communicate the spatial goal turned out to be a straight line that connects the starting point of the robots to their goal. The implementation of these movements is described in Appendix A.1.

However, the number of robots we tested for were 5 and 15, and we chose these numbers so that they were distinguishable between themselves and also different from the number used in the previous works [9, 10], where the experiments were ran on teams of 20 robots, thus allowing




Fig. 8. Point of view of the user (height approximately 1.7 m). It replicates the proximal interaction scenario in which the user is working alongside the robots.

us to evaluate the effect of the number of robots on the correct recognition of the coordination objectives.

In contrast, the dependent variables of the experiments are the correctness of the answers and the time to answer taken in both parts of the experiments. The term *dependent variables* refers to the quantities that we measure during the experiments and upon which we will carry out the statistical analysis. We consider these dependent variables, because they represent the effectiveness of a legible system, namely, we define the ability to communicate in Definition 1 in terms of providing the correct information and in a short time.

4.4 Factorial Design

To investigate the statistical correlation between the independent and the dependent variables, we built a factorial plan [49]. Given the independent variables of the two parts of the experiment, as reported in Section 4.3, we had 4×2 and 3×2 combinations in the two parts, respectively. In addition, following the design of experiment methodology [49], we took some precautions to reduce the disturbing factors that might influence the answer of the experiments. It is worth noting that both the environment and some precautions were inspired by our previous work [9, 10].

4.4.1 Point of View. The position of the user into the environment was fixed throughout the experiments to avoid additional disturbing factors. Figure 8 reports the height of the camera view into the Unity scene, chosen to be approximately the standard height of a human being, simulating the condition of a proximal interaction in which the user is standing close to the robots. Furthermore, the motion of the user's head was mapped into a coherent motion of the point of view, in such a way to reduce the motion sickness [33] that could incur when the movement of the head does not match with a concordant movement of what the user sees.

4.4.2 Position of the Goals. In the experimental scenario, described in Section 4.1, five goals were equally distributed in front of the user, as shown in Figure 9. Users were explicitly informed of the fact that all the goals had equal probability of being visited. The two extreme goals, located at the far left and far right of the setup, were never used in the experiment to ensure that all the analyzed goals had two neighbors.

The goals were randomized and the same combination of goal and trial could not happen twice for a given user. Thus, in the second part of the experiment the goal used in the first part for a given combination of coordination objective and number of robots was excluded.

4.4.3 Learning Effect. To reduce the effects of the learning effect that takes place as participants advance through the user study, we introduced some additional parts to the experimental

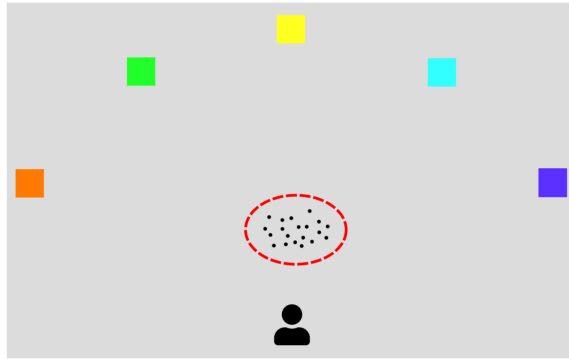


Fig. 9. Top-down view of the position of the goals (*colored cubes*), of the robots at the beginning of the experiments (*circled in red*), and of the user (*user icon*). The scope is to simulate a proximal interaction scenario, where the user shared the same environment with the robots.

procedure described in Section 4.2. In particular, the experiments started with a small tutorial to familiarize the user with the input method. Then a practice scenario allowed the users to test the input method and to get used to the surrounding environment. In the second part of the experiment, the users were able to see again the demonstrative examples as many times as they wanted. In addition, the order of the trials of the factorial plan were randomized.

4.4.4 Movement of the Robots. The controllers of the robots, described in Section 3, were implemented in MATLAB, and the trajectories of all the robots were saved offline to be replicated exactly in the Unity environment during the experiments. This affords high repeatability during the experiments, which results in a reduction of the disturbing factor that may be introduced from random modification of the trajectories in a real-time implementation.

All the trajectories were normalized in terms of total duration, which was fixed to 60 s: Namely, the group of robots reached the spatial goal in the given time. We experimentally found out that this is a reasonable time (given also the previous experiences [9, 10]). The initial positions of the robots were equal for all the trials and were chosen to reach the specified formation as quickly as possible. This was done to balance out the formation objective with the consensus and coverage objectives in terms of time required for the multi-robot team to achieve a distinguishable behavior, since we discovered that, for the set of parameters chosen in the experiments (see Appendix A.4), the formation was the slowest to converge to a recognizable pattern.

4.5 Users

We performed the experiments on a group of 21 users (age 25.76 ± 5.37 y.o., 2 females and 19 males). All the participants were students or researchers within the engineering department, but none of them were involved in the project nor they partook in the building of the experiment. Regarding their robotics background, six users had a multi-robot background and five did not have any robotic knowledge, namely, they were completely new to the subject. All the users signed a consent form and they were instructed on the experimental procedure as described in Section 4.2. We conducted a within-subjects study (see Reference [57] for details) and, according to this methodology, each user tests all the possible combinations of the independent variables in both parts of the experiment. This choice allowed us to obtain a sufficient amount of data from a limited number of participants. The main drawback of this method may be the learning effect or the fatigue, which we addressed with the precautions described in Section 4.4.

Last, the possible disorientation that may arise through the use of a virtual reality headset was counterbalanced by the fact that the majority of the users claimed to have tried at least once virtual reality (10 users had used it once or twice, 7 more than twice, and 4 users had never used it before).

4.6 Statistical Tools

To process the data obtained from the experiments, we utilized standard statistical tools [49], which will be detailed hereafter.

Usually, the first step in a statistical analysis is to check the significance of the independent variables over the dependent variables. The null hypothesis is that there is no statistical relevance, which can be rejected or supported based on the calculation of the p -value. Considering a significance level of 0.05 implies that if the p -value is smaller than 0.05, then we can reject the null hypothesis with a 5% of risk of asserting a wrong assumption. To calculate the p -value, two different approaches are used depending on the kind of dependent variable: categorical, for variables like blood type or binary answers, or continuous, for variables such as time or temperature. In particular, for the analysis of our user study, we have one categorical variable (correctness of the answer given by the users) and one continuous variable (time taken to give the answer).

For categorical variables, the significance tests build upon a contingency table, which is a table where the rows are the levels of the independent variable and the columns are the categories of the dependent variable. The cells are filled based on the result of each trial (i.e., based on the value of the independent variable and of the value of the dependent variable, we add an entry to the corresponding cell). Consequently, each cell has a number that corresponds to the frequency of that given combination of independent and dependent variable. Once the contingency table is built, a significance test can be carried out, with the Pearson's Chi-Square test being a typical choice to determine the significance. However, Fisher's exact test is used instead when one frequency of the contingency table is less than five. Both tests measure the independence between the column and row variables of the given contingency table.

For continuous variables the standard tool is ANOVA (Analysis of variance), which explains whether the independent variables are statistically relevant with respect to the dependent variable. Additionally, a post-hoc test, such as the Tukey-Kramer multiple comparison procedure, can be carried out. This test brings out the differences among the groups (i.e., the combination of the levels of the independent variables) and it can highlight if the statistical relevance is caused by a particular group or not.

5 RESULTS

This section outlines the statistical analysis and discussion of the data acquired during the experiments. The objective is to understand whether the multi-robot system investigated in the experiments is legible, in terms of Definition 1, by analyzing the relationship between the independent variables and legibility.

5.1 Spatial Goal

Along the lines of Definition 1, the spatial goal of the multi-robot system is the first information that the system should be able to communicate to be legible. To study the statistical relevance of the independent variables over the ability of the multi-robot system to communicate its spatial goal, we identify the following null hypotheses:

H1—Null hypothesis: *There is no statistical relationship between the independent variables and the legibility of the system, in terms of correctness for the understanding of the spatial goal.*

Table 1. Contingency Table of Correctness of Answers with Respect to Coordination Objective

Coordination objective \ Correctness	Correctness		
	Wrong	Right	Total
Consensus	0	42	42
Coverage	1	41	42
Formation	0	42	42
Coverage time-varying	1	41	42
Total	2	166	168

Table 2. Contingency Table of Correctness of Answers with Respect to the Number of Robots

Number of robots \ Correctness	Correctness		
	Wrong	Right	Total
5	0	84	84
15	2	82	84
Total	2	166	168

Table 3. p -Values from the Fisher's Exact Test on Correctness

Independent variable	p -value
Coordination objective	0.6856
Number of robots	0.4970

H2—Null hypothesis: *There is no statistical relationship between the independent variables and the legibility of the system, in terms of response time for the understanding of the spatial goal.*

For analyzing the correctness of the goal, we built the corresponding contingency tables, Tables 1 and 2. Since correctness is a categorical variable, we need two separate tables, one for each independent variable. Given that the results are not suitable for Pearson's Chi-Square test (see Section 4.6 for details), we instead carried out a Fisher's exact test. The p -value for both results was high (see Table 3), as expected from the contingency tables. In fact, most of the answers were correct, independently from the coordination objective or the number of robots (no evident differences among the rows of the contingency tables considering the same column). Actually, the percentage of success 98.81% (166 right answers over 168 total) would bring to a lack of relevance of any independent variable. This result is consistent with the outcome of the previous works [9, 10], because it seems that the main parameter that affects the communication of the spatial goal is the trajectory of the multi-robot system, and in this experiment all the trials had a straight trajectory toward their goal, independently from the coordination objectives or the number of robots. As the Fisher's exact tests have not highlighted any statistical relevance, we did not further investigate this relationship, and we cannot reject **H1**. This result is also confirmed by the survey, in which the answer to the question *Do you think you have correctly understood the spatial goal of the robot?* was 4.238 ± 0.625 .

As regards the response time, given its continuous nature, we used ANOVA to discover if the variables chosen are statistically relevant. We discarded the trials with a wrong answer, because a penalty function would have to be considered, which can interfere with the data. Table 4 reports

Table 4. p -Values from the ANOVA on Time Response

Independent variable	p -value
Coordination objective	0.0309
Number of robots	0.0811
Coordination objective-Number of robots	<0.0001




Fig. 10. Results of the post-hoc test. The trial formation with five robots, which is the one significantly different from the other trials, is highlighted in red. The probable cause for this disparity is the time needed to achieve the formation, which affects the time performance of this coordination objective.

the p -values for the independent variables considered in the ANOVA, from which we can deduce that coordination objective and its interaction with the number of robots are relevant (in the tables the relevant variables are highlighted in grey). Instead, the number of robots is not relevant. However it is worth noting that the Tukey-Kramer post-hoc test highlighted that these results were mostly influenced by the trial with formation and five robots, which turned out to be the hardest to understand. This is reported in Figure 10. As a result, we reject **H2** for the coordination objective and its interaction with the number of robots.

5.1.1 Influence of Coordination Objective on Spatial Goal Identification. The results of Section 5.1 about **H1** were verified also with the data of the second part of the experiment, even if one of the coordination objectives was not considered (time-varying coverage). The percentage of success was 97.6% (123 correct answers in understanding the spatial goal over 126 total), and all the tests for the response time achieved the same results. To verify that the users were not influenced by the fact of being asked to answer to two queries simultaneously, we carried out an additional ANOVA adding the part of the experiment as independent variable. In this analysis, we neglected the trials with the time-varying coverage of the first part to avoid an unbalanced study. The null hypothesis for this study is:

H3—Null hypothesis: *There is no statistical relationship between the two parts of the experiment and the legibility of the system, in terms of response time for the understanding of the spatial goal, namely, the addition of the second answer did not affect the legibility.*

This additional analysis confirms **H3** (p -value 0.468 in the ANOVA). Hence, in the second part, the users were not influenced by the additional question on the coordination objective. The survey

confirms also this aspect: the users answered 1.714 ± 0.956 to the question *In the second part, do you think that choosing both the goal and the coordination objective distracted you?*

5.1.2 Coverage Control. In the first part of the experiment, we implemented both the standard and the time-varying version of coverage control. To understand if there is some difference between the two versions, we evaluate the following hypothesis:

H4—Null hypothesis: *There is no statistical relationship between the type of coverage control and the legibility of the system, in terms of response time for the understanding of the spatial goal.*

To investigate **H4**, we carried out ANOVA only between the two types of coverage to understand whether the time-varying aspect can affect the legibility. From the ANOVA, no statistical relevance emerged (p -value 0.0874), but considering the mean (standard coverage 13.98 ± 6.65 s, time-varying coverage 16.69 ± 5.9 s), we can assert that using the standard coverage for the second part is legitimate, since standard coverage is slightly more legible. In fact, the time-varying coverage looks similar to the consensus, because the robots converge to a point and then they move toward a goal, i.e., the density function moves toward the goal and, as a consequence, the robots. Therefore, we cannot reject **H4**.

5.2 Coordination Objective

For the second ability of Definition 1, we study the ability of the multi-robot system to communicate its coordination objective. The null hypotheses are:

H5—Null hypothesis: *There is no statistical relationship between the independent variables and the legibility of the system, in terms of correctness for the understanding of the coordination objective.*

H6—Null hypothesis: *There is no statistical relationship between the independent variables and the legibility of the system, in terms of response time for the understanding of the coordination objective.*

For the analysis of the correctness of the recognition of the coordination objective in the second part of the experiment, we proceeded as for the correctness of the goal. Tables 5 and 6 report the contingency tables for the coordination objective and for the number of robots, respectively. Given that the entries are sufficiently numerous, we carried out the Pearson's Chi-Square test, and we discovered that both independent variables are significant (Table 7 reports the p -values). Hence, we can reject **H5**: specifically, this result shows that both the coordination objective and the different number of robots influence the ability of the user in understanding the coordination objective. It is worth noting that even the users perceived a good confidence in recognizing the coordination objective: the users answered 3.667 ± 0.796 to the question *Were you able to distinguish among the coordination objectives?* In addition, the real percentage of correctness of this part was 68% (86 correct answers over 126 total), which means that the users were able to recognize the coordination objectives.

We removed the wrong answers to proceed with the analysis of the response time, as we did for testing **H2** in Section 5.1. We checked the balance of the remaining data, namely, if some combination had too many or too few entries. Table 8 reports the p -values calculated with ANOVA, which allows us to reject **H6**. However, as already pointed out for the response time for the goal recognition in Section 5.1, the post-hoc test reported that the trial with formation and five robots is the one that differs most from the others, namely, it was the most difficult to recognize.

About the survey of this last part, we report that consensus was considered as the easiest coordination objective to recognize (100% of the users chose consensus in the question *Which coordination objective was the easiest to recognize?*). After the experiment, many users asserted that

Table 5. Contingency Table of Correctness of Answers with Respect to Coordination Objective in the Recognition of the Coordination Objective

Coordination objective \ Correctness	Wrong	Right	Total
	Consensus	5	37
Formation	21	21	42
Coverage	14	18	42
Total	40	86	126

Table 6. Contingency Table of Correctness of Answers with Respect to the Number of Robots in the Recognition of the Coordination Objective

Number of robots \ Correctness	Wrong	Right	Total
	5	27	36
15	13	50	63
Total	40	86	126

Table 7. p -Value From the Pearson's Chi-square on Correctness for the Coordination Objective Recognition

Independent variable	p -value
Coordination objective	<0.0001
Number of robots	0.0074

Table 8. p -Value from ANOVA on Response Time for the Coordination Objective Recognition

Independent variable	p -value
Coordination objective	<0.0001
Number of robots	0.0154
Coordination objective-Number of robots	0.0045

consensus was the simplest, because they associated the small occupied area to that coordination objective, and, on the other side, formation and coverage occupied a similar area. Instead, coverage had been identified as the hardest to recognize (71.4% of the users answered coverage to the question *Which coordination objective was the hardest to recognize?* and the remaining answered formation, confirming even more the preference on consensus).

Formation might seem as the easiest coordination objective to recognize, but it turned out to be the most difficult in this user study. Based on comments received by the users after performing the experiments, we inferred that this is due probably to the time needed to create the formation and the high dispersion of the robots when they perform this coordination objective. In addition the point of view is very different from the point of view with which formation control is usually presented, namely, top-view (see Figure 11). In fact, with the position of the camera inside the virtual environment (as reported in Figure 8) and sharing the environment it is difficult to recognize any shape. Differently, consensus was easy to recognize, because the spatial occupancy is easy to identify.




Fig. 11. Issue on formation recognition caused by the point of view of the user. On the left, with the standard top-view it is easy to recognize and to understand the coordination objective. Instead, on the right, the view that the users had during the experiment: the formation is hard to recognize due to the viewing angle.

6 IMPLICATIONS AND LIMITATIONS

The results collected from the user study offer significant insights into the legibility of multi-robot systems, as will be discussed in detail in Section 7. However, the main assumptions upon which the user study was design lead to a number of limitations, detailed hereafter.

6.1 Legibility Definition

One of the key contributions of this article is in the formal definition of legibility for a multi-robot system. Such definition is then exploited to derive the metric for the evaluation of the user study results. While this definition is given in a generic manner, suitable for diversified case studies, some application scenarios may not be well represented. For instance, the spatial goal is defined as a single constant point in space, while some application scenarios may require, for instance, periodic trajectories. Regarding the coordination objectives, we considered their definition in terms of coordinated control strategies, while other approaches (e.g., planning methodologies [64, 74]) might be also be interesting to investigate.

6.2 Virtual Reality Simulations and Lack of Realism

The experiments were performed in virtual reality only. This allowed us to precisely control the experimental conditions, reducing (as much as possible) the presence of external, uncontrolled, confounding factors. At the same time, the lack of realism of the virtual reality simulation environment reduces the generality of the results, which could be enriched considering the disturbances and variability that would arise from real-world experiments.

As a design choice, the user was not allowed to move within the virtual reality environment. Considering the user’s motion would be another possible confounding factor that could lead to interesting insights. In fact, variation in the viewpoint of the user is known to be a relevant factor in the context of communication between humans and robots [2, 26].

6.3 Design Choices

The collected results are dependent on the design of the spatial targets, as well as on the design of the coordination objectives. In particular, while the selected coordination objectives are recognized, in the literature, as prototypical examples, different coordinated control strategies might be considered as well: in this case, one could expect to obtain additional insights, from a user study. At the same time, the specific parameter design of each coordination objective (e.g., the shape of the formation, or the number of robots) might also influence the results. Investigating the effect of

the choice of parameters for each control action on the legibility would represent an interesting line for future research.

6.4 Reduced Diversity in the User Study Participants

The user study was conducted recruiting students and researchers within the engineering department: as such, diversity in terms of gender and background knowledge was not guaranteed. Considering these additional dimensions for future user studies would increase the generality of the results.

7 DISCUSSION AND CONCLUSIONS

In this article, we investigated the legibility of multi-robot systems. We extended the definition of legibility for this kind of systems, incorporating both the ability of communicating a spatial goal and a coordination objective. Then, we carried out a user study to answer to the question “Can a multi-robot system be legible, considering Definition 1?” In the experiment, we investigated both the legibility with respect to the spatial goals and the coordination objectives. The user study was conducted in a virtual reality setup, and the acquired data were analyzed with standard statistical tools.

About the ability of communicating the spatial goal, i.e., the first part of Definition 1, we can merge the information gathered in this article and in the previous works [9, 10]. From the results of the performed user studies, we can infer that, at least for simple tasks (as the ones considered in this article, where spatial goals are sufficiently diversified), an overall straight trajectory is preferable, together with a limited inter-robot distance.

Large inter-robot distances could be due by high dispersion caused by strong artificial potential fields, or by a big variance of the Gaussian probability density function of the coverage control, or again by the initial procedure to create the shape in formation control. The performed user studies allow us to infer the conclusion that (at least for the conditions considered in this article) a high level of cohesiveness of the group along the path toward the goal helps in making the overall motion more legible for the human observer.

Regarding the communication of a coordination objective, i.e., the second part of Definition 1, it is important to highlight the characteristics of the desired behavior. In fact, for consensus its inherent characteristic of agglomeration was easily identifiable. Thus, consensus was the easiest coordination objective to recognize. Instead, in the case of formation control, the major feature is the geometry of the formation. If this geometry does not emerge quickly enough, then the identification of the coordination objective is significantly hindered. To improve the legibility of this coordination objective, we would probably need to create an understandable formation even with the different point of view typical of proximal interaction. Further investigations are needed to understand the relationship between legibility and the shape of the formation. In addition, working toward the legibility of multi-robot systems executing complex behaviors will imply balancing the legibility aspect with the performance of the overall system.

From the results of the user study, we can establish that multi-robot systems are legible according to Definition 1. In fact, multi-robot systems can communicate effectively their spatial goal and their coordination objective. While these results are satisfactory, we are still witnessing the dawn of the study of legibility of multi-robot systems and a lot of work has to be done to improve the methods to communicate and to merge legible motion with other tasks of multi-robot systems.

Regarding future directions, we plan to perform a comparative study to assess, in a quantitative manner, what coordination objectives are more easily confused with what other ones, to provide better guidelines in the choice of parameters for the control laws. Furthermore, we plan to explore if explicit communication outperforms implicit communication, and by how much. At the moment,

we cannot state that legible motion is better nor worst than explicit communication both in terms of efficiency or perceived integration by the user. However, we think that implicit communication should be integrated with standard communication means, as it represents an essential part of everyday communication among human beings. In this context, this article advances the understanding of legibility for multi-robot systems, pushing toward the integration of these systems with humans.

We remark that the environment where the robot motion took place was intentionally left as simple as possible, to focus the study on the motion characteristics only. As a future line of research, we plan to study the interplay between the motion patterns and the environment design, in terms of legibility: As suggested in Reference [37], in fact, the design of the environment plays a crucial role in this regard.

Another interesting line of future work is represented by multi-robot systems interacting with multiple users. While works can be found in the literature that consider these aspects for the single robot case (see, e.g., Reference [23]), we plan to extend this context to a multi-robot scenario.

APPENDIX

A EXPERIMENTAL IMPLEMENTATION

This article evaluates the legibility of multi-robot systems via a user study where participants identify key aspects of the system's behavior. To this end, as described in Section 4, a series of virtual reality scenarios are shown to the participants for them to identify where the multi-robot team is headed (spatial objective) and how the team is self-organized (coordination objective). This appendix expands on the control laws outlined in Section 3, providing a detailed explanation of the controller implementation displayed in the user study.

A.1 Spatial Drive of the Multi-robot Team

A core legibility aspect studied in this article relates to the inference of the spatial goal of the multi-robot system, i.e., where is the team headed toward within the domain. In the consensus or the formation control protocols, one way of displacing the multi-robot team throughout the domain is to independently command one of the robots in the team, a *leader*, to move toward a certain location. As a consequence of the inter-agent interactions encoded in the control law Equations (1) and (2), the movement of this leader robot influences the positioning of the whole team in the domain. If we denote as $x_L \in \mathbb{R}^2$ the position of the leader robot, then we can make the consensus and formation coordination objectives move toward a spatial goal by executing the following go-to-goal controller,

$$\dot{x}_L = \kappa_{goal}(x_{goal} - x_L),$$

where $x_{goal} \in \mathbb{R}^2$ denotes the desired location of the team, and $\kappa_{goal} \in \mathbb{R}_+$ is a proportional gain, respectively. With the purpose of minimally disturbing the motions corresponding to the multi-robot coordination objectives, we discard selecting the leader among of the N robots displayed in the experiment and, rather, pick the leader robot to be virtual.

Coverage, however, naturally affords the spatial allocation of the multi-robot team within the domain. Under the protocol presented in Equations (3)–(5), the robots are displaced toward the spatial goals by specifying the density function, ϕ , in Equation (3) to be maximal around the goal and decrease as the distance to the goal increments. To this end, we choose the density functions to be bivariate unimodal Gaussians of the form

$$\phi(q) = \frac{k}{\sigma^2} \exp\left(-\frac{\|q - \mu\|^2}{2\sigma^2}\right), \quad (7)$$

where $q \in D$ is a point in the domain, $k \in \mathbb{R}_+$ is a proportional constant, $\sigma \in \mathbb{R}_+$ is the variance, and $\mu \in \mathbb{R}^2$ is the mean of the function. By making the mean, μ , coincide with the spatial goal, x_{goal} , the team moves toward the objective location when executing the control law in Equation (4). The parameters chosen for the density function are specified in Appendix A.4.

In the case of the time-varying coverage in Section 3.3.1, we consider a density function $\phi(x, t)$ of the form in Equation (7), where the dependency on t stems from considering a time-varying mean, $\mu(t)$. The evolution of the mean serves the purpose of guiding the multi-robot team from its initial location toward the spatial goal. Consequently, μ coincides with the mean position of the robots at the start and moves on a straight line toward the spatial goal as the experiment progresses according to

$$\mu(t) = \bar{x}(t_0) + \frac{x_{goal} - \bar{x}(t_0)}{t_F - t_0}(t - t_0),$$

where $\bar{x} \in \mathbb{R}^2$ denotes the mean position of the robots, $\bar{x}(t) = \sum_{i \in \mathcal{N}} x_i(t)/N$, and t_0 and t_F the initial and final time, respectively. For the user study, $t_0 = 0$ s and $t_F = 60$ s.

A.2 Collision Avoidance

The controllers presented in Section 3 model robots as points that move in the space according to single integrator dynamics. However, in reality, robots have a finite footprint and mass, and collisions must be avoided in general. We used **Control Barrier Functions (CBFs)** as in Reference [70] to prevent collisions among the robots. This choice was motivated by the minimal perturbation that the CBFs introduce on the coordinated behavior of the robots as opposed to other methods such as repulsive potential fields, which can introduce inter-robots connections that modify the overall behavior of the system (as seen in References [9, 10]) and potentially hinder the legibility of the coordinated behavior.

Let $d_{min} \in \mathbb{R}_+$ be the minimum distance between any pair of robots with positions x_i and x_j , we can define the following CBF to ensure pairwise collision-avoidance,

$$h_{ij}(x_i, x_j) = \|x_i - x_j\|^2 - d_{min}^2.$$

A safe velocity $u^* \in \mathbb{R}^{2N}$ for the multi-robot team can be found solving the following quadratic program (see Reference [72] for further details),

$$\begin{aligned} u^*(x) &= \operatorname{argmin}_{u \in \mathbb{R}^{2N}} \|u - u^{nom}\|^2 \\ \text{s.t. } &\frac{\partial h_{ij}(x_i, x_j)}{\partial x_i} u \geq -\gamma h_{ij}(x_i, x_j)^3, \\ &\forall i \in \{1, \dots, N-1\}, j \in \{i+1, \dots, N\}, \end{aligned}$$

where $u^{nom} \in \mathbb{R}^{2N}$ is the nominal input obtained from the controllers in Section 3 and $\gamma \in \mathbb{R}_+$. The controller u^* avoids collisions while remaining as close as possible to the nominal controller.

A.3 Formation Control

To execute the formation controller in Equation (2), we need a formation Δ specifying the desired pairwise distances between robots. As is well known, regular polygonal shapes are easy to recognize, for users [17, 24]. Hence, to make the result of the formation control clearly distinguishable from the other coordination objectives, for the user study in Section 4, we choose formations where the robots are arranged in concentric regular polygons. In particular, for $N = 5$, we choose a pentagon and for $N = 15$, two concentric polygons (one heptagon and one octagon). For the achieved formation see Figures 12(a) and 12(b), respectively.




Fig. 12. Formation for the different number of robots. These formation are obtained applying the formation control law (2), without any spatial drive, to clarify the achieved shape. The blue lines correspond to distances, δ_{ij} , to be maintained between robots.

Table 9. Parameters for the Coordination Strategies

Parameter	Consensus	Formation	Coverage
κ	1	1.5	1
κ_{goal}	$1.5 * \kappa * N$	$0.5 * \kappa$	-
\dot{x} saturation [m/s]	0.4	0.4	0.4
d_{min} [m]	0.1	0.1	0.1
σ	-	-	1

The formations are generated as follows: let d_g be the distance between goals (they are set equally spaced out), then the formation is circumscribed by a circle of radius $R_{out} = d_g/3$. For the case of $N = 15$, the inner polygon is concentric to the outer octagon and circumscribed by a polygon of radius $R_{in} = R_{out}/4$. With the geometries defined, the formation Δ is obtained by placing the robots on the vertices of the polygons. The pairwise distances δ_{ij} for a robot i will correspond to the Euclidean distance between its corresponding vertex and the rest of the robots in the polygon. An example of how these distances are generated is shown in Figure 12(a).

A.4 Choice of Parameters

The parameters used for each coordination objective are reported in Table 9. The saturation on the velocity, \dot{x} saturation, was chosen to produce a realistic behavior of the robots, which have a maximum speed at which they can move. The coverage parameter σ was chosen to produce a similar spatial occupancy, in terms of area occupied by the convex hull of the positions of the robots, as the formation control (see Section A.3 for further details). However, consensus has a smaller dispersion given its inherent operation and the agglomeration of the robots was only affected by the safety distance, d_{min} . The parameters corresponding to the proportional gain κ , which properly scales the input, and κ_{goal} were set to make the team achieve the spatial goal within the specified time frame of 60 s.

A.5 Examples of trajectories

The following figures depict examples of trajectories, represented from top-down, followed by the robots during the experiments. In the figures, red asterisks represent the initial position of the robots, green circles represent intermediate positions, and blue triangles represent the final positions.

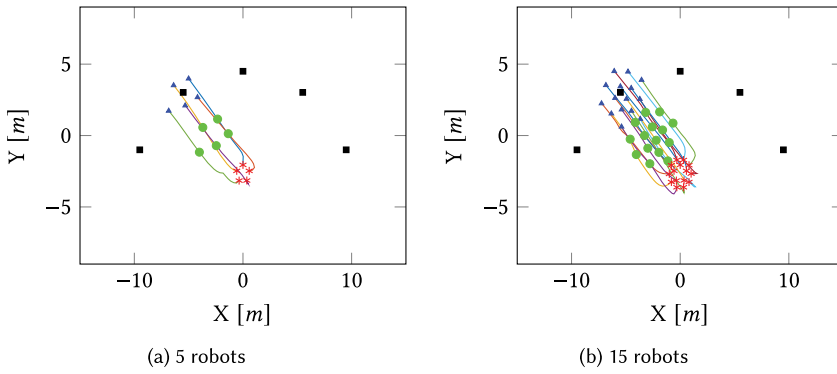


Fig. 13. Example of top-down trajectories toward the left top goal, for the time-varying coverage control coordination objective. The red asterisks represent the initial position of the robots, the green circle the intermediate ones, and the blue triangle the final ones.




Fig. 14. Example of top-down trajectories toward the left top goal, for the coverage control coordination objective. The red asterisks represent the initial position of the robots, the green circle the intermediate ones, and the blue triangle the final ones.




Fig. 15. Example of top-down trajectories toward the left top goal, for the consensus coordination objective. The red asterisks represent the initial position of the robots, the green circle the intermediate ones, and the blue triangle the final ones.




Fig. 16. Example of top-down trajectories toward the left top goal, for the formation control coordination objective. The red asterisks represent the initial position of the robots, the green circle the intermediate ones, and the blue triangle the final ones.

Table 10. First Part of the Experiment: Time to Provide the Answer, Based on the Spatial Goal

Goal	Mean and std
G0	16.5211 ± 6.8270 s
G1	13.5277 ± 5.4885 s
G2	17.9285 ± 7.0081 s

A.6 Quantitative Results

We report hereafter the quantitative results recorded in the user studies: for both parts of the experiments, mean and standard deviation of the response time are reported.

The three spatial goals are identified, respectively, as G0, G1, G2.

Each coordination behavior was executed with 5 and 15 robots, and the respective combinations are identified as follows:

- COVtv5: Time-varying coverage control, 5 robots (executed in the first part only);
- COVtv15: Time-varying coverage control, 15 robots (executed in the first part only);
- COV5: Coverage control, 5 robots;
- COV15: Coverage control, 15 robots;
- CON5: Consensus, 5 robots;
- CON15: Consensus, 15 robots;
- F5: Formation control, 5 robots;
- F15: Formation control, 15 robots.

We remind the readers that, in the second part of the experiment, the time-varying coverage was not utilized.

Tables 10, 11, and 12 collect the values (mean and standard deviation) of the time taken to provide the answer in the first part of the experiments, considering, respectively, only the spatial goal, only the coordination objective and the number of robots, and both aspects together.

Tables 13, 14, and 15 collect the values (mean and standard deviation) of the time taken to provide the answer on the spatial goal, in the second part of the experiments, considering, respectively,

Table 11. First Part of the Experiment: Time to Provide the Answer, Based on the Coordination Objective and Number of Robots

Combination	Mean and std
COVtv5	15.7137 ± 5.5406 s
COVtv15	16.8091 ± 6.3161 s
COV5	11.7760 ± 6.1830 s
COV15	15.8666 ± 6.6326 s
CON5	13.5878 ± 4.9804 s
CON15	14.9078 ± 3.9684 s
F5	25.9041 ± 5.3961 s
F15	13.0891 ± 3.4690 s

Table 12. First Part of the Experiment: Time to Provide the Answer, Considering the Spatial Goal, Combined with Coordination Objective and Number of Robots (Mean and Standard Deviation)

Goal	Coordination objective and number of robots							
	COVtv5	COVtv15	COV5	COV15	CON5	CON15	F5	F15
G0	15.4466 ± 5.5441	17.6065 ± 6.2949	12.3671 ± 7.3302	17.2483 ± 10.3072	16.2340 ± 4.6033	18.9170 ± 0.9037	31.4565 ± 1.0373	—
G1	—	11.3840 ± 7.7075	11.7355 ± 5.7812	13.8306 ± 5.7178	11.5966 ± 4.0080	13.9054 ± 3.9488	22.4144 ± 0.8037	—
G2	21.0560 ± 0	15.4565 ± 5.5784	7.5330 ± 0	21.9565 ± 3.1741	18.1046 ± 4.8619	15.2839 ± 4.0334	26.3573 ± 5.9049	13.0891 ± 3.4690

Table 13. Second Part of the Experiment: Time to Provide the Answer on the Spatial Goal, Based on the Spatial Goal

Goal	Mean and std
G0	16.1587 ± 6.7637 s
G1	13.2096 ± 5.6655 s
G2	16.8527 ± 6.7756 s

Table 14. Second Part of the Experiment: Time to Provide the Answer on the Spatial Goal, Based on the Coordination Objective and Number of Robots

Combination	Mean and std
COV5	12.5859 ± 5.8023 s
COV15	15.9054 ± 6.2474 s
CON5	12.1532 ± 3.6131 s
CON15	12.7797 ± 4.0241 s
F5	24.7832 ± 5.1234 s
F15	13.2513 ± 3.7694 s

only the spatial goal, only the coordination objective and the number of robots, and both aspects together.

Tables 16, 17, and 18 collect the values (mean and standard deviation) of the time taken to provide the answer on the coordination objective, in the second part of the experiments, considering, respectively, only the spatial goal, only the coordination objective and the number of robots, and both aspects together.

Table 15. Second Part of the Experiment: Time to Provide the Answer on the Spatial Goal, Considering the Spatial Goal, Combined with Coordination Objective and Number of Robots (Mean and Standard Deviation)

Goal	Coordination objective and number of robots					
	COV5	COV15	CON5	CON15	F5	F15
G0	14.3789 ± 5.1508	16.2796 ± 6.8830	12.5622 ± 3.3390	13.5449 ± 3.6765	29.2770 ± 2.9017	17.8680 ± 4.5600
G1	8.0634 ± 3.1209	12.4337 ± 2.1464	9.3547 ± 2.1777	9.8978 ± 1.5452	22.9968 ± 1.6696	12.1651 ± 2.6980
G2	16.5637 ± 7.8605	17.1992 ± 6.0746	13.6225 ± 4.8630	13.8156 ± 4.8536	22.9692 ± 7.6283	–

Table 16. Second Part of the Experiment: Time to Provide the Answer on the Coordination Objective, Based on the Spatial Goal

Goal	Mean and std
G0	12.3250 ± 8.1746 s
G1	11.5157 ± 5.8187 s
G2	12.8764 ± 7.7876 s

Table 17. Second Part of the Experiment: Time to Provide the Answer on the Coordination Objective, Based on the Coordination Objective and Number of Robots

Combination	Mean and std
COV5	11.0853 ± 5.2705 s
COV15	13.9896 ± 5.7466 s
CON5	8.2434 ± 5.8996 s
CON15	7.3398 ± 4.9104 s
F5	19.8253 ± 9.6131 s
F15	12.3888 ± 3.6836 s

Table 18. Second Part of the Experiment: Time to Provide the Answer on the Coordination Objective, Considering the Spatial Goal, Combined with Coordination Objective and Number of Robots (Mean and Standard Deviation)

Goal	Coordination objective and number of robots					
	COV5	COV15	CON5	CON15	F5	F15
G0	12.7619 ± 6.3283	13.4790 ± 6.5305	6.6565 ± 4.9344	8.2044 ± 6.4868	26.3880 ± 7.6682	12.6540 ± 5.1092
G1	8.4301 ± 2.7616	14.6710 ± 1.7589	7.4522 ± 3.5358	6.6712 ± 5.0501	15.8297 ± 8.9930	12.3264 ± 3.4705
G2	11.1333 ± 3.8120	15.2655 ± 5.4251	14.1920 ± 7.9505	6.8931 ± 3.2921	19.2560 ± 10.0318	–

ACKNOWLEDGMENT

Authors thank Noemi Canovi for her help in the user study.

REFERENCES

- [1] H. Ando, Y. Oasa, I. Suzuki, and M. Yamashita. 1999. Distributed memoryless point convergence algorithm for mobile robots with limited visibility. *IEEE Trans. Robot. Autom.* 15, 5 (1999), 818–828. <https://doi.org/10.1109/70.795787>
- [2] Christopher Bodden, Daniel Rakita, Bilge Mutlu, and Michael Gleicher. 2016. Evaluating intent-expressive robot arm motion. In *Proceedings of the 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'16)*. IEEE, 658–663.

- [3] Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. 2013. Swarm robotics: A review from the swarm engineering perspective. *Swarm Intell.* 7, 1 (Mar. 2013), 1–41.
- [4] Cynthia Breazeal. 2003. Toward sociable robots. *Robot. Auton. Syst.* 42, 3 (2003), 167–175.
- [5] Beatrice Capelli and Lorenzo Sabattini. 2019. Human-friendly multi-robot systems: Legibility analysis. In *Proceedings of the International Workshop on Human-Friendly Robotics (HFR'19)*. Springer, 15–28.
- [6] Beatrice Capelli, Maria Santos, and Lorenzo Sabattini. 2023. Legibility of Multi-robot Systems: Examples of coordination objectives. (May 2023). <https://doi.org/10.5281/zenodo.7896482>
- [7] Beatrice Capelli, Maria Santos, and Lorenzo Sabattini. 2023. Legibility of Multi-Robot Systems: Virtual Reality experimental setup. (May 2023). <https://doi.org/10.5281/zenodo.7896403>
- [8] Beatrice Capelli, Maria Santos, and Lorenzo Sabattini. 2024. Legibility of Multi-Robot Systems: virtual reality experiments. (Feb. 2024). <https://doi.org/10.5281/zenodo.10611702>
- [9] Beatrice Capelli, Cristian Secchi, and Lorenzo Sabattini. 2019. Communication through motion: legibility of multi-robot systems. In *Proceedings of the International Symposium on Multi-Robot and Multi-Agent Systems (MRS'19)*. IEEE, 126–132.
- [10] Beatrice Capelli, Valeria Villani, Cristian Secchi, and Lorenzo Sabattini. 2019. Understanding multi-robot systems: On the concept of legibility. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'19)*. IEEE, 7355–7361.
- [11] Tathagata Chakraborti, Anagha Kulkarni, Sarath Sreedharan, David E. Smith, and Subbarao Kambhampati. 2019. Expliability? Legibility? Predictability? Transparency? Privacy? Security? The emerging landscape of interpretable agent behavior. In *Proceedings of the 29th International Conference on Automated Planning and Scheduling*, Vol. 29. 86–96.
- [12] Tathagata Chakraborti, Sarath Sreedharan, Anagha Kulkarni, and Subbarao Kambhampati. 2018. Projection-aware task planning and execution for human-in-the-loop operation of robots in a mixed-reality workspace. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'18)*. IEEE, 4476–4482.
- [13] Jorge Cortés and Magnus Egerstedt. 2017. Coordinated control of multi-robot systems: A survey. *SICE J. Control, Measure. Syst. Integr.* 10, 6 (2017), 495–503.
- [14] J. Cortes, S. Martinez, and F. Bullo. 2006. Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. *IEEE Trans. Autom. Control* 51, 8 (2006), 1289–1298. <https://doi.org/10.1109/TAC.2006.878713>
- [15] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. 2004. Coverage control for mobile sensing networks. *IEEE Trans. Robot. Autom.* 20, 2 (2004), 243–255.
- [16] Maartje M. A. De Graaf and Bertram F. Malle. 2019. People's explanations of robot behavior subtly reveal mental state inferences. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI'19)*. IEEE, 239–248.
- [17] J. A. Deutsch. 1955. A theory of shape recognition. *Brit. J. Psychol.* 46, 1 (1955), 30.
- [18] Griffin Dietz, Jane L. E., Peter Washington, Lawrence H. Kim, and Sean Follmer. 2017. Human perception of swarm robot motion. In *Proceedings of the CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 2520–2527.
- [19] Anca Dragan and Siddhartha Srinivasa. 2013. Generating legible motion. In *Proceedings of the 9th Annual Conference on Robotics: Science and Systems (RSS'13)*.
- [20] Anca D. Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S. Srinivasa. 2015. Effects of robot motion on human-robot collaboration. In *Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI'15)*. IEEE, 51–58.
- [21] Anca D. Dragan, Kenton C. T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *Proceedings of the 8th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI'13)*. IEEE Press, 301–308.
- [22] Qiang Du, Vance Faber, and Max Gunzburger. 1999. Centroidal voronoi tessellations: Applications and algorithms. *SIAM Rev.* 41, 4 (1999), 637–676. <https://doi.org/10.1137/S0036144599352836>
- [23] Miguel Faria, Francisco S Melo, and Ana Paiva. 2021. Understanding robots: Making robots more legible in multi-party interactions. In *Proceedings of the 30th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN'21)*. IEEE, 1031–1036.
- [24] Yang Feng, Qiong Wu, Keisuke Okamoto, Jiajia Yang, Satoshi Takahashi, Yoshimichi Ejima, and Jinglong Wu. 2017. A basic study on regular polygons recognition of central and peripheral vision field for virtual reality. In *Proceedings of the IEEE International Conference on Mechatronics and Automation (ICMA'17)*. IEEE, 1738–1743.
- [25] Federica Ferraguti, Chiara Talignani Landi, Silvia Costi, Marcello Bonfè, Saverio Farsoni, Cristian Secchi, and Cesare Fantuzzi. 2020. Safety barrier functions and multi-camera tracking for human–robot shared environment. *Robot. Auton. Syst.* 124 (2020), 103388.
- [26] Paul Fletcher, Angeline Luther, Brittany Duncan, and Carrick Detweiler. 2021. Investigation of unmanned aerial vehicle gesture perceptibility and impact of viewpoint variance. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'21)*. IEEE, 3531–3537.

- [27] Terrence Fong, Illah Nourbakhsh, and Kerstin Dautenhahn. 2003. A survey of socially interactive robots. *Robot. Auton. Syst.* 42, 3-4 (2003), 143–166.
- [28] Antonio Franchi, Paolo Robuffo Giordano, Cristian Secchi, Hyoung Il Son, and Heinrich H. Bühlhoff. 2011. A passivity-based decentralized approach for the bilateral teleoperation of a group of UAVs with switching topology. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'11)*. IEEE, 898–905.
- [29] Fabrizio Ghiringhelli, Jérôme Guzzi, Gianni A. Di Caro, Vincenzo Caglioti, Luca M. Gambardella, and Alessandro Giusti. 2014. Interactive augmented reality for understanding and analyzing multi-robot systems. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'14)*. IEEE, 1195–1201.
- [30] Gerd Hirzinger, Klaus Landzettel, Bernhard Brunner, Max Fischer, Carsten Preusche, Detlef Reintsema, Alin Albu-Schäffer, Günter Schreiber, and Bernhard-Michael Steinmetz. 2004. DLR's robotics technologies for on-orbit servicing. *Adv. Robot.* 18, 2 (2004), 139–174.
- [31] Meng Ji and Magnus Egerstedt. 2007. Distributed coordination control of multiagent systems while preserving connectedness. *IEEE Trans. Robot.* 23, 4 (2007), 693–703. <https://doi.org/10.1109/TRO.2007.900638>
- [32] George K. Karavas and Panagiotis Artemiadis. 2015. On the effect of swarm collective behavior on human perception: Towards brain-swarm interfaces. In *Proceedings of the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI'15)*. IEEE, 172–177.
- [33] Behrang Keshavarz, Heiko Hecht, and Ben D. Lawson. 2014. Visually induced motion sickness: Causes, characteristics, and countermeasures. In *Handbook of Virtual Environments*. CRC Press, 652–703.
- [34] Lawrence H. Kim and Sean Follmer. 2021. Generating legible and glanceable swarm robot motion through trajectory, collective behavior, and pre-attentive processing features. *ACM Trans. Human-Robot Interact.* 10, 3 (2021), 1–25.
- [35] John Klingner, Anshul Kanakia, Nicholas Farrow, Dustin Reishus, and Nikolaus Correll. 2014. A stick-slip omnidirectional powertrain for low-cost swarm robotics: Mechanism, calibration, and control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. 846–851. <https://doi.org/10.1109/IROS.2014.6942658>
- [36] Andreas Kolling, Phillip Walker, Nilanjan Chakraborty, Katia Sycara, and Michael Lewis. 2016. Human interaction with robot swarms: A survey. *IEEE Trans. Hum.-Mach. Syst.* 46, 1 (2016), 9–26.
- [37] Anagha Kulkarni, Sarath Sreedharan, Sarah Keren, Tathagata Chakraborti, David E. Smith, and Subbarao Kambhampati. 2020. Designing environments conducive to interpretable robot behavior. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'20)*. IEEE, 10982–10989.
- [38] Anagha Kulkarni, Siddharth Srivastava, and Subbarao Kambhampati. 2019. A unified framework for planning in adversarial and cooperative environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 2479–2487.
- [39] Przemyslaw A. Lasota, Terrence Fong, Julie A. Shah et al. 2017. A survey of methods for safe human-robot interaction. *Found. Trends Robot.* 5, 4 (2017), 261–349.
- [40] Sung G. Lee, Yancy Diaz-Mercado, and Magnus Egerstedt. 2015. Multirobot control using time-varying density functions. *IEEE Trans. Robot.* 31, 2 (2015), 489–493.
- [41] N. E. Leonard and E. Fiorelli. 2001. Virtual leaders, artificial potentials and coordinated control of groups. In *Proceedings of the 40th IEEE Conference on Decision and Control*, Vol. 3. 2968–2973 vol.3. <https://doi.org/10.1109/CDC.2001.980728>
- [42] F. Levillain, D. St-Onge, E. Zibetti, and G. Beltrame. 2018. More than the sum of its parts: Assessing the coherence and expressivity of a robotic swarm. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'18)*. 583–588. <https://doi.org/10.1109/ROMAN.2018.8525640>
- [43] J. Lin, A. S. Morse, and B. D. O. Anderson. 2007. The multi-agent rendezvous problem. part 1: The synchronous case. *SIAM J. Control Optimiz.* 46, 6 (Dec. 2007), 2096–2119. <https://doi.org/10.1137/040620552>
- [44] S. Lloyd. 1982. Least squares quantization in PCM. *IEEE Trans. Info. Theory* 28, 2 (1982), 129–137. <https://doi.org/10.1109/TIT.1982.1056489>
- [45] Alyxander David May, Christian Dondrup, and Marc Hanheide. 2015. Show me your moves! Conveying navigation intention of a mobile robot to humans. In *Proceedings of the European Conference on Mobile Robots (ECMR'15)*. IEEE, 1–6.
- [46] Mehran Mesbahi and Magnus Egerstedt. 2010. *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press.
- [47] Marco Minelli, Federica Ferraguti, Nicola Piccinelli, Riccardo Muradore, and Cristian Secchi. 2019. An energy-shared two-layer approach for multi-master-multi-slave bilateral teleoperation systems. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'19)*. IEEE, 423–429.
- [48] Francesco Mondada, Michael Bonani, Xavier Raemy, James Pugh, Christopher Cianci, Adam Klaptoc, Stéphane Magnenat, Jean christophe Zufferey, Dario Floreano, and Alcherio Martinoli. 2009. The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*. 59–65.
- [49] Douglas C. Montgomery. 2017. *Design and Analysis of Experiments*. John Wiley & Sons.
- [50] Robin R. Murphy. 2004. Trial by fire [rescue robots]. *IEEE Robot. Autom. Mag.* 11, 3 (2004), 50–61.

- [51] Jawad Nagi, Alessandro Giusti, Luca M. Gambardella, and Gianni A. Di Caro. 2014. Human-swarm interaction using spatial gestures. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'14)*. IEEE, 3834–3841.
- [52] Günter Niemeyer, Carsten Preusche, Stefano Stramigioli, and Dongjun Lee. 2016. Telerobotics. In *Springer Handbook of Robotics*. Springer, 1085–1108.
- [53] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. 2000. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley.
- [54] R. Olfati-Saber. 2006. Flocking for multi-agent dynamic systems: algorithms and theory. *IEEE Trans. Autom. Control* 51, 3 (2006), 401–420. <https://doi.org/10.1109/TAC.2005.864190>
- [55] Frank Papenmeier, Meike Uhrig, and Alexandra Kirsch. 2019. Human understanding of robot motion: The role of velocity and orientation. *Int. J. Soc. Robot.* 11, 1 (2019), 75–88.
- [56] Jayam Patel, Yicong Xu, and Carlo Pinciroli. 2019. Mixed-granularity human-swarm interaction. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'19)*. IEEE, 1059–1065.
- [57] Neil J. Salkind. 2010. *Encyclopedia of Research Design*. Vol. 1. Sage.
- [58] Maria Santos and Magnus Egerstedt. 2020. From motions to emotions: Can the fundamental emotions be expressed in a robot swarm? *Int. J. Soc. Robot.* (2020), 1–14.
- [59] Jean Scholtz, Mary Theofanos, and Brian Antonishek. 2006. Development of a test bed for evaluating human-robot performance for explosive ordnance disposal robots. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*. 10–17.
- [60] Melanie Schranz, Martina Umlauf, Micha Sende, and Wilfried Elmenreich. 2020. Swarm robotic behaviors and current applications. *Front. Robot. AI* 7 (2020), 36.
- [61] Adriane E. Seiffert, Sean Timothy Hayes, Caroline E. Harriott, and Julie A. Adams. 2015. Motion perception of biological swarms. In *Proceedings of the 37th Annual Meeting of the Cognitive Science Society (CogSci'15)*.
- [62] Aaron St. Clair and Maja Mataric. 2015. How robot verbal feedback can improve team performance in human-robot task collaborations. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI'15)*. ACM, 213–220.
- [63] David St-Onge, Ulysse Côté-Allard, Kyrre Glette, Benoit Gosselin, and Giovanni Beltrame. 2019. Engaging with robotic swarms: Commands from expressive motion. *ACM Trans. Hum.-Robot Interact.* 8, 2 (2019), 1–26.
- [64] Michal Štolba, Antonín Komenda, and Daniel Kovacs. 2016. Competition of distributed and multiagent planners (codmap). In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30.
- [65] Daniel Szafir, Bilge Mutlu, and Terrence Fong. 2014. Communication of intent in assistive free flyers. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI'14)*. ACM, 358–365.
- [66] Daniel Szafir, Bilge Mutlu, and Terry Fong. 2015. Communicating directionality in flying robots. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI'15)*. ACM, 19–26.
- [67] Valeria Villani, Beatrice Capelli, and Lorenzo Sabattini. 2018. Use of virtual reality for the evaluation of human-robot interaction systems in complex scenarios. In *Proceedings of the 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'18)*. IEEE, 422–427.
- [68] Glenn Wagner and Howie Choset. 2015. Gaussian reconstruction of swarm behavior from partial data. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'15)*. IEEE, 5864–5870.
- [69] Phillip Walker, Michael Lewis, and Katia Sycara. 2016. Characterizing human perception of emergent swarm behaviors. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC'16)*. IEEE, 002436–002441.
- [70] Li Wang, Aaron D. Ames, and Magnus Egerstedt. 2017. Safety barrier certificates for collisions-free multirobot systems. *IEEE Trans. Robot.* 33, 3 (2017), 661–674. <https://doi.org/10.1109/TRO.2017.2659727>
- [71] Sean Wilson, Ruben Gámeros, Michael Sheely, Matthew Lin, Kathryn Dover, Robert Gevorkyan, Matt Haberland, Andrea Bertozzi, and Spring Berman. 2016. Pheeno, A versatile swarm robotic research and education platform. *IEEE Robot. Autom. Lett.* 1, 2 (2016), 884–891. <https://doi.org/10.1109/LRA.2016.2524987>
- [72] Sean Wilson, Paul Glotfelter, Li Wang, Siddharth Mayya, Gennaro Notomista, Mark Mote, and Magnus Egerstedt. 2020. The robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems. *IEEE Control Syst. Mag.* 40, 1 (2020), 26–44. <https://doi.org/10.1109/MCS.2019.2949973>
- [73] Changbin Yu, Brian D. O. Anderson, Soura Dasgupta, and Barış Fidan. 2009. Control of minimally persistent formations in the plane. *SIAM J. Control Optimiz.* 48, 1 (2009), 206–233. <https://doi.org/10.1137/060678592>
- [74] Yu Zhang, Sarath Sreedharan, and Subbarao Kambhampati. 2016. A formal analysis of required cooperation in multi-agent planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, Vol. 26. 335–343.
- [75] Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, Hankz Hankui Zhuo, and Subbarao Kambhampati. 2017. Plan explicability and predictability for robot task planning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'17)*. IEEE, 1313–1320.

Received 25 July 2022; revised 17 October 2023; accepted 25 January 2024