HuGoS: A Multi-user Virtual Environment for Studying Human–Human Swarm Intelligence

Nicolas Coucke[1,2(✉)], Mary Katherine Heinrich[1], Axel Cleeremans[2], and Marco Dorigo[1]

[1] IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{nicolas.coucke,mary.katherine.heinrich,axcleer,mdorigo}@ulb.ac.be
[2] CO3, Center for Research in Cognition and Neurosciences, Université Libre de Bruxelles, Brussels, Belgium

Abstract. The research topic of human–human swam intelligence includes many mechanisms that need to be studied in controlled experiment conditions with multiple human subjects. Virtual environments are a useful tool to isolate specific human interactions for study, but current platforms support only a small scope of possible research areas. In this paper, we present HuGoS—'Humans Go Swarming'—a multi-user virtual environment in Unity, as a comprehensive tool for experimentation in human–human swarm intelligence. We identify possible experiment classes for studying human collective behavior, and equip our virtual environment with sufficient features to support each of these experiment classes. We then demonstrate the functionality of the virtual environment in simple examples for three of the experiment classes: human collective decision making, human social learning strategies, and agent-level human interaction with artificial swarms, including robot swarms.

1 Introduction

Human–human swarm intelligence is a broad field of study [20], including topics such as crowd dynamics [31], online social networks [22], and collective problem solving [39]. While some studies of human group behavior use data collection from real-world systems, such as social networks [33], many study types require controlled experiment conditions. As self-organization in human groups normally occurs within the context of other mechanisms and influences, a comprehensive tool for studying human–human swarm intelligence must enable the experimenter to artificially limit human capabilities of perception and communication, according to the given experiment. Virtual environments have been proposed as tools to isolate and study specific aspects of human interaction [1].

In this paper, we develop a virtual environment for experiments with multiple human subjects. To be comprehensive, the environment needs to support studies in three main topics of human–human swarm intelligence. First, humans often use simple mechanisms and strict self-organization to coordinate (e.g., in

M. Dorigo et al. (Eds.): ANTS 2020, LNCS 12421, pp. 161–175, 2020.
https://doi.org/10.1007/978-3-030-60376-2_13

human crowds [31]), displaying behaviors similar to those observed in artificial swarms and animal groups. Second, humans also use more complex mechanisms (e.g., advanced negotiation, or hierarchical social structures) that are often formed via self-organization. For instance, hierarchy can be self-organized according to response speeds of individuals [21], or strengths of preexisting interpersonal ties [7]. Third, comparative studies between human groups and artificial swarms are relevant even for complex mechanisms, as self-organized leadership and hierarchy have recently been studied not only in humans, but also in groups of non-human animals [13] and groups of robots [24]. In this paper, we propose HuGoS—'Humans Go Swarming'—a multi-user virtual environment that supports research and experimentation in each of these three topics. In HuGoS, human participants interact via avatars in a controlled experiment setup, capable of supporting both simple and complex interactions. HuGoS also supports avatars controlled by artificial agents, enabling comparative studies between human and artificial behaviors.

1.1 Related Work

A number of multi-user virtual environments have been developed for studies with human groups, in two main categories. The first category of environments use collective human gameplay or other interactions as tools for solving computationally intensive problems [2,6,9,18,23,44], rather than studying underlying cognitive or behavioral mechanisms. The second category of virtual environments are those developed primarily to study the mechanisms of collective human behavior. For collective decision making, the *UNUM* platform [40,41], also referred to as Swarm AI®, supports a group of participants that collaboratively explore a decision space [26]. For physical coordination, a Unity implementation supports human-like avatars with first-person view for the study of crowd behaviors [32,45,53]. Finally, a third category supports the study of leadership—specifically, the impact of better informed individuals on implicit leadership (e.g., the *HoneyComb* game for human crowd movement [4]). These existing environments are mostly developed for a specific task. For instance, in the *UNUM* platform, each player controls a 'magnet' that exerts influence on a 'puck.' This platform could not be easily re-purposed to investigate, for instance, a task involving environment exploration. To our knowledge, there are no existing platforms that are versatile enough to be used for a wide variety of topics in human–human swarm intelligence. In this paper, we target the contribution of a platform that can comprehensively cover this field of study. For instance, in addition to the implicit leadership studied in [4], a comprehensive platform should also be able to study explicit leadership, such as 'follower' functionalities developed in online trading networks [19]. Furthermore, the existing platforms log experiment data, such as positions of avatars in the virtual environment, but conduct analysis externally. We target a platform in which recorded data can also be analyzed internally, allowing real-time feedback to be incorporated in the experiment. This capability enables the study of, for instance, the relationship between group behavior and different types of performance feedback.

Existing tools used for robot simulation support 3D environments in which artificial agents can interact both with each other and with the environment. ARGoS [37] is a tool built specifically for robot swarms, while other tools such as ROS [38] or Webots [27] are built for robots generally. These versatile tools can be adapted to a wide variety of experiment scenarios, and many types of data collection and analysis. However, they have limited applicability to the study of human behavior. A few studies have looked at human–swarm interaction using general tools for robots (e.g., using ROS [49] or Webots [48]). In these setups, humans are able to give high-level directions to individual robots. However, the humans cannot act independently of the robots, as they cannot control their own avatars with first-person view. Also, the approaches do not demonstrate multiple human users in one environment.

2 Design of HuGoS: 'Humans Go Swarming'

We target the design of a multi-user virtual environment—HuGoS: 'Humans Go Swarming'—that can be used as a tool to study human collective behavior generally, including collective decision-making, collaborative task performance, and the emergence of leadership. HuGoS should also support the study of differences and similarities between human swarm intelligence and artificial swarm intelligence, and the interactions between human and artificial agents.

2.1 Experimentation Scope for Human–Human Swarm Intelligence

There are several classes of experiments that HuGoS must support, to facilitate comprehensive study of human–human swarm intelligence. We base this experimentation scope on existing studies of robot and artificial swarms, such that these behaviors could also be studied in humans. The first class of experiments is physical coordination between individuals, as in flocking and self-assembly (e.g., [42]). In HuGoS, this would require a minimal environment in which participants control avatars, whose positions are continuously recorded for analysis. The second class involves observation of environmental features. In cooperative navigation, for example, agents might extract and share information to find the shortest path in an environment (e.g., [8]). In best-of-n decision-making, a swarm might choose the best of several options based on observations of the environment (e.g., [46]). In HuGoS, this requires that the environment be populated with game objects that act as obstacles, or represent environment features with observable discrete or continuous properties. In order to study decision making based on external information, of the type studied in human groups in platforms such as *UNUM* [40], landmarks in the environment could be labelled with each option, and participants could use their avatar positions in relation to the landmarks to indicate their opinions. The third class involves the agents making changes to the environment, such as in *stigmergic* communication (e.g., [12]) or in the performance of a task such as collective construction (e.g., [51]). In HuGoS, this requires game objects that can be manipulated or modified by avatars—for instance, immovable environment features with modifiable properties such as color, or movable objects such as construction blocks. All classes require HuGoS to enable various types of direct and indirect communication between players, in ways that can be expressly limited. For example, avatars' view capabilities may be limited such that players can see only their immediate neighbors, not all avatars. Or, in a task such as collective construction, players may be able to see only the construction blocks, not the other avatars. In all classes, studies might include a comparison between human and artificial behaviors, or collaboration between human and artificial agents. This requires, in addition to human-controlled avatars, that HuGoS supports artificial agents whose avatars may be indistinguishable for human players. As it might be fruitful to replicate these experiments in real setups, we also need to integrate robot models into HuGoS (such as those developed for the ARGoS multi-robot simulator [37]).

2.2 Features of the HuGoS Virtual Environment

Unity 3D Game Engine. HuGoS is built in Unity, a 3D game development platform that can support intelligent agents in a physically realistic game environment [14]. In Unity, basic building blocks of virtual environments are termed *game objects*. Each game object represents a physical 3D object within the game environment that is subject to physics engines and is linked to specific C# back-end scripts. Through these scripts, each game object in HuGoS can be: i) passive, ii) controlled by simple rule-based behaviors but immobile, iii) mobile and equipped with a controller to act as an artificial agent, or iv) mobile and controlled by a human player. We refer to game objects in HuGoS as *avatars* if they act as artificial agents or are controlled by human players. We refer to game objects as *landmarks* if they are immobile, whether passive or controlled by simple rule-based behaviors (e.g., change display color to match that of neighboring landmark). Using Unity's networking capabilities, we organize the multi-user architecture of HuGoS as follows.

HuGoS initiates on a server, and each new player joins as client on that server. Throughout the experiments, the server logs all data recorded from HuGoS. The activities taking place on the server are divided into three modules: the player module, the environment module, and the task module (Fig. 1(a)). The environment module tracks landmarks, including changes made to them by either players or controllers, and passes those updates to the clients. The player module tracks player actions, and mediates any communication between clients. The task module tracks and analyzes the progress of the specific experiment, which can optionally be shared with player clients.

Avatar Capabilities. Each player controls an avatar that is situated in the virtual environment. The capabilities of the avatars in a given experiment setup are defined in the player module. In HuGoS, players have a third-person view of their avatar through a virtual camera that follows the avatar position and rotation. Players move their avatar by pressing four user-customizable keys (e.g.,

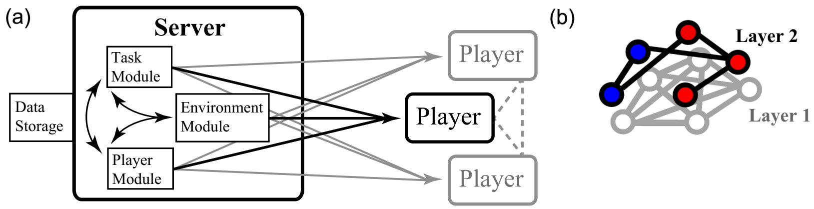

Fig. 1. (a) Program architecture. The server contains the task, player, and environment modules. Communication with player clients is managed by modules (dark arrows represent communication between modules and player; lighter arrows illustrate that same communication to multiple players; dashed lines represent communication between players, as mediated by the player module). (b) Example of layered player networks. Layer 1 is fully connected (visibility of other avatars), while layer 2 is partly connected (visibility of the other avatars' color opinions).

WSAD), and rotate via the left/right arrow keys or cursor movement. Depending on the experiment scenario, specific additional actions can be activated for the avatars. For example, the player can be permitted to manipulate the environment by clicking on landmarks to grab them, then moving and releasing the cursor to move them. Indirect communication between players can occur via changes to the environment, for instance by moving landmarks, or by changes to display features of the player's avatar, such as color. Direct communication can also be permitted—and limited as desired—by sending symbols, or written or spoken messages. In the player module of HuGoS, the players' environment perception is controlled firstly by changing the field of view (FOV) of the player. A player that has a limited top-down avatar FOV (Fig. 3(c)) can only perceive the environment in a small perimeter, while a player that has an oblique avatar FOV (Fig. 3(a)) can see a much greater proportion of the environment, in the viewed direction. Unlimited top-down FOV is also possible, giving a player global view (Fig. 3(b)). Additionally, avatars and landmarks can be programmed to be invisible to players, or to be visible only for a subset of players.

Player interactions can be modulated by changing the structure of the *player networks* in the player module, which are directed graphs. If a player network is fully connected, for instance, then every player can interact with all other players in the way associated to that network. Player networks manages different types of interaction, and have independently defined structures. For example, a fully connected network might be defined for viewing avatar positions, while a sparsely connected network might be defined for viewing avatar colors (Fig. 1(b)). Player networks also govern explicit message passing between players. As connections are directed (e.g., player 1 might be able to see player 2, while player 2 cannot see player 1), the information privileges of players can be made hierarchical. Certain players can have higher node indegrees or outdegrees. The structure of player networks can be changed during experiment runtime, and can optionally be triggered by the players. For instance, players might be permitted to 'follow' another player by clicking on its avatar, causing their own decisions to

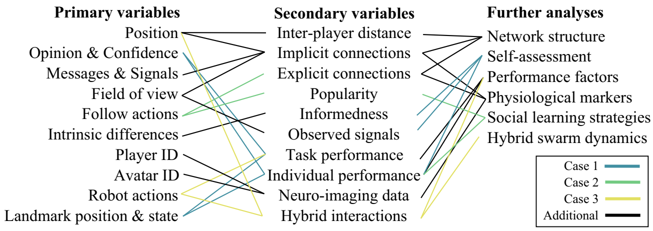

Fig. 2. Primary variables from the environment can be used for analysis specific to the experimental conditions, to calculate secondary variables and conduct further analysis.

automatically copy those of the followed player, until that player is unfollowed (see Sect. 3.2). The ability to control communication links between players also allows for comparison between limited communication networks and fully connected communication networks. This can facilitate the study of information cascades, bias in the group, or dysfunctional dynamics that may lead to low performance.

Data Types (Primary Variables). Data about the players, environment, and task are logged for analysis. Each player has a unique anonymized player ID, and each avatar has an avatar ID. These two IDs are important in cases where players switch avatar identities between trials, so that the behaviors of specific players can be analyzed separately from the features accumulated by a shared avatar. Additionally, the player IDs are important in experiment setups that involve players' physical environments (e.g., players occupy the same physical environment and can communicate, or players' physiological data is monitored, such as EEG). Avatar capabilities, positions, FOVs, and actions are all logged, according to avatar ID. These logs enable calculation of other simple data about avatars, such as which other avatars are in one avatar's FOV. Messages passed by avatars are also logged, including the content, time, sender avatar ID, and receiver avatar ID. All other player interactions are also tracked and logged as events—for instance, a player choosing to follow another player—again including content, time, and sender and receiver IDs. Changes in the environment are also logged, including positions and states of landmarks. When artificial agents such as robots are included in a setup, any data specific to that agent and experiment is logged. For instance, in a setup with models of e-puck robots [29], proximity sensing and motor control might be logged.

Analysis Types (Secondary Variables). The data logged as primary variables (i.e., recorded directly) allow many secondary variables to be calculated and analyzed during runtime or during post-processing (Fig. 2). Here, we use task performance as an illustrative example. Task performance can be continuously calculated by the task module, according to the specific scenario. In a flocking scenario, the task performance would depend on player positions; in decision-making, on player opinions. Once task performance is calculated, additional analysis might assess, for instance, how this performance relates to the in-game behavior of players. Player behavior in this case might be represented by distances between avatars, the network of implicit connections between individuals that occur when avatars enter each other's FOVs, or the network of direct messages between players with connection weights representing message frequency. The primary variables also allow analysis of individual behavior, that can be used to give feedback to players during the experiment. For instance, in a collective decision-making scenario, comparing individual opinion to overall task performance yields the relative player performance. If this is provided as feedback to players, players can use it to determine and display their opinion confidence. If the calculated player performance is not provided to the player when the player determines opinion confidence, then a comparison of these two variables will yield the player's *self-assessment* (i.e., the ability to evaluate their own performance). Using player IDs, out-of-game data can also be used in post-analysis. For example, each player might be asked to fill in a questionnaire about personality traits or subjective experience during the game. In an extended out-of-game setup, gameplay could even be linked to real-time physiological recordings, such as eye-tracking, ECG or EDA tracking of stress, or neural recordings via EEG or fMRI. Such extensions could be used to analyze the connection between individual cognitive mechanisms and collective performance during gameplay.

3 Demonstration of HuGoS Features

We demonstrate the suitability of HuGoS for the defined experimentation scope in three case studies. First, we demonstrate human players in a basic scenario previously studied in robot swarms [47]. Second, we demonstrate an aspect of human behavior that is outside typical swarm intelligence studies—specifically, the establishment of leader–follower relationships, similar to behavior mimicking in human trading networks [19]. Third, we demonstrate artificial agent avatars in a basic swarm intelligence scenario, and demonstrate interaction between human-controlled avatars and artificial agent avatars (specifically, robot avatars).

3.1 Case 1: Collective Decision Making

Collective decision making is widely studied in artificial swarms (e.g., [41,47]). We implement a setup based on that of [47] in HuGoS, but with consensus to be reached by human players rather than kilobots (Fig. 3). The task is to reach a consensus about the predominant color in the environment, which is populated with cylindrical landmarks that are randomly distributed and colored red or blue (Fig. 3(b)). The difficulty of the task can be adjusted by changing the color ratio and the density of landmarks. In this example there are 1 000 landmarks, in a color ratio of 55%-45%. At initiation of the run, the task module randomly determines whether red or blue will be the most prevalent. Each player controls an avatar that initiates at a random position. Half of the avatars initiate with

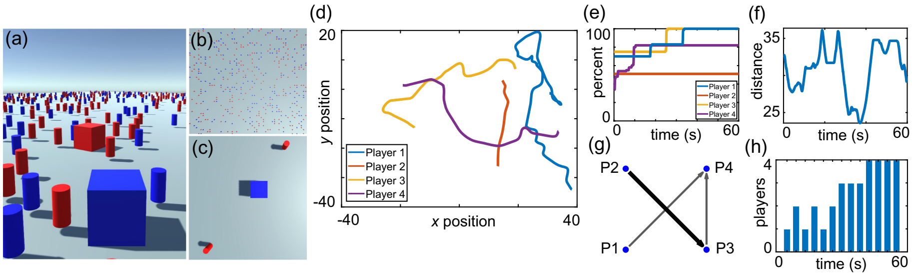

Fig. 3. (a) Player view on display monitor. The player controls an avatar (blue cube in the center of the screen) and can see the avatar of another player (red cube), in addition to cylindrical landmarks. (b) Top view of the environment with random distribution of blue and red cylinders. (c) Player view with limited information—views avatar from the top and can see only local information. (d–h) Trial of collective decision making with four human-controlled avatars. (d) Trajectories traveled by the avatars. (e) Percentage of the environment seen by each player. (f) Average Euclidean distance between avatars. (g) Network of player–player view (time 0–60 s); connection weights indicate total view time. (h) Number of avatars displaying the correct color opinion. (Color figure online)

the opinion red, and the other half blue, assigned randomly. The players can switch their color opinions by pressing a key, and can move and rotate without restriction. The player has a third-person avatar view through a virtual camera that follows the avatar (Fig. 3(a)). By moving and rotating the avatar, and thus the FOV, a player can explore the environment from different perspectives and make a subjective estimate of the majority color in the playing field. A player can see the movements and color changes of all avatars in the FOV. On the server, the player module passes avatar colors to the task module, where task performance is calculated, according to the homogeneity of avatar opinions, and whether the majority opinion matches the dominant color in the environment. Gameplay ends when all players have the same color, or when a time limit is reached. Variables are logged and calculated during runtime, including the task performance (Fig. 3(h)), avatar FOV, and avatar positions (Fig. 3(d)). The average euclidean distance between every two avatars is calculated (Fig. 3(f)), as is the percentage of cylinders cumulatively observed in the FOV (Fig. 3(e)), and the directed graph of all avatars' appearances in others' FOVs (Fig. 3(g)).

3.2 Case 2: Social Learning Strategies

In swarm intelligence in artificial agents, interactions are typically not chosen by individuals, but rather happen indiscriminately through random encounters. In humans and certain other animals, selective interaction can have a substantial impact on collective behavior, and has been widely studied in the context of *social learning* (cf. [17,52]). Humans in social learning scenarios use this selectivity to choose when, from whom, and what to learn [17]. In a collective decision making scenario, the dynamics of this selectivity would have a significant impact on the

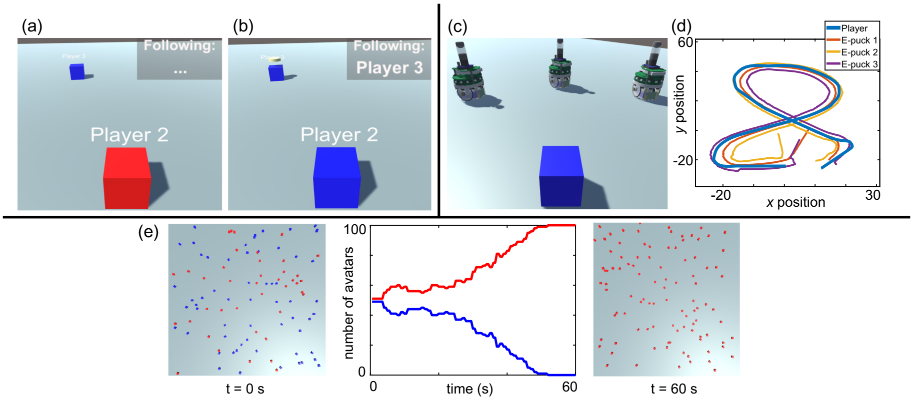

Fig. 4. (a–b) In Case 2: Explicit leader–follower relationships, in the view of player 2. (a) Player 2 acting individually. (b) Player 2 following player 3. (c–d) In Case 3: Human–artificial interaction. (c) Human player-controlled avatar (blue cube) followed by artificial agent avatars (e-puck robots). (d) Trajectories of a human player (in blue) being followed by three e-puck robots. (e) In Case 3: 100 artificial agent avatars performing random walk and using majority rule for color opinions.

outcome of collective behavior. Player strategies for selective interaction could be studied implicitly after calculating the information seen in the avatar FOV, as in Figs. 3(e,g). Strategies for selective interaction could be studied more explicitly, via a function for the self-organization of explicit leader–follower relationships between players. Individuals might choose, for example, to follow the most prestigious individual (i.e., who already has most followers) [5,17], or the most successful individual (i.e., highest individual performance, for the task) [11,17]. We therefore add events for leader–follower relationships to HuGoS. During gameplay, players choose either to act for themselves (Fig. 4(a)), or to 'follow' another avatar and copy its actions (Fig. 4(b)). A player can choose to follow another by clicking on that player's avatar. Once the relationship is established, the following status is displayed to the follower player in a dialogue box, and a tiara appears above the leader avatar in the follower's display (Fig. 4(b)). While the relationship exists, the follower player no longer is in control of the avatar, which automatically copies the behavior of the leader avatar. In a scenario similar to Case 1, this would mean that the follower's motion copies that of the leader, and the follower automatically adopts the color of the leader. A follower can at any time decide to follow a different leader by clicking on the corresponding avatar, or decide to act independently again by clicking on the ground plane.

3.3 Case 3: Interaction with Artificial Agents

HuGoS also includes mobile avatars that are artificial agents, rather than being controlled by human players, as well as immobile landmarks equipped with rule-based behaviors for their display features. This enables straightforward comparison of artificial swarms and human collective behavior in the same virtual environment. Analogous to case study 1, we implement a simple collective decision making scenario with artificial agent avatars, where each avatar initiates as either blue or red. At each simulation step, the avatars move via a random walk, and update their color opinion using a simple majority rule [30]. That is, an avatar updates its color to the color opinion held by the majority of avatars in a 5 m radius (with the total environment size being $80 \times 80\,\mathrm{m}^2$). We show the results of this behavior in a swarm of 100 artificial agent avatars—Fig. 4(e) gives the top view at initiation, top view after consensus, and the percentage of color opinions in the avatar swarm over time. Artificial agents also allow the study of hybrid human–robot avatar swarms, as human players can interact with simulated robots. To demonstrate this, we transfer to Unity a Pi-puck [28] robot model, based on a model developed for the multi-physics multi-robot simulator ARGoS [37]. As Unity includes built-in physics engines, existing ARGoS kinematics and dynamics robot models could also be transferred into HuGoS, including those that have already been calibrated to real hardware (e.g., [36]). To demonstrate interaction between players and robots, we programmed the simulated robots to detect a nearby human player avatar, and to move in the direction of that avatar when the distance becomes too large (Fig. 4(c)). The robot avatars also sense their distance to other robots and move away from each other if they get too close. We show the trajectories of all avatars in a setup with one human player avatar and three robot avatars (Fig. 4(d)).

4 Discussion

We have introduced a novel multi-player virtual environment that is suited for studying human behavior in swarm intelligence scenarios. HuGoS logs primary variables related to actions of individual players and multiple players, and variables that depend on the environment. These primary variables, recorded directly, are then used to calculate secondary variables (e.g., the collective performance of avatars for a given task).

Many research topics proposed here are typically studied with in-person experimental setups in which participants interact directly. Such setups allow for many types of interaction that are not possible in HuGoS, such as eye contact, speech characteristics, and body language. We do not propose that HuGoS is a replacement for in-person studies. Rather, our objective is to use the simplifications of a virtual environment for complementary studies that provide new insights by isolating certain aspects of previously studied dynamics. Isolating specific aspects can be useful in studying human behavior [1], for instance by helping to disentangle different forms of interaction that would be too closely related in an in-person setup.

As HuGoS targets human players, some open questions remain. As pointed out in [3], it is possible that human players will pay less attention to other individuals when there are less communication channels available or when too many other individuals are present. Individual behavior is also highly dependent on whether or not players are convinced that other avatars are actually human-controlled [3]. Another challenge with human players is to keep players engaged, and motivated to perform well in the task. Performance trackers such as leaderboards might motivate players [50]. However, in many cases, these external rewards do not increase players' intrinsic motivation [25], and in some cases have been shown to not improve performance [35]. For players to be optimally engaged—that is, making the game intrinsically motivating—there should be a clear goal, players should feel in control of the outcome and receive regular feedback on their actions, and the task should be neither too easy nor too difficult [15,34]. Another problem might be reproducibility. Humans have a wide range of individual behaviours, and, the more varieties of behavior possible in a task, the less likely it might be that players would follow similar behavior over several trials. Players might also start with widely different prior skills on a task; some players might be more familiar with video games than others. Also, personality differences might play a role in the game. For example, players that are more socially dominant might be less inclined to follow other players. These variations can be assessed with questionnaires prior to the experiment [16]. Human participants that are recruited in one experiment setting are also likely to have a mostly homogeneous cultural background, which has been shown to have a substantial impact on behavior [10]. Interesting changes in behaviors might be explored when players from multiple backgrounds collaborate in a game. Finally, recruiting participants might be a challenge. If experiments are conducted in one shared computer client room, there is the advantage of control and overview of participant behavior (cf. [53]). If we alternatively run HuGoS as an online browser game, it might make it easier to recruit large numbers of participants, but can make participant behavior less controllable [6,43].

5 Conclusions

We have designed and presented HuGoS, a multi-user virtual environment that supports the study of human–human interactions and group behaviors relevant to the topic of swarm intelligence. HuGoS supports collection of the primary data types required to analyze relevant aspects of human behavior. The environment's flexibility allows for implementation of a wide variety of swarm intelligence scenarios. We have demonstrated three simple cases of such scenarios, demonstrating support for: 1) studying human behavior in tasks typically studied in artificial swarms, such as best-of-n collective decision making; 2) studying new behaviors that may be especially relevant to human collective behavior, such as the self-organization of hierarchical social structures; and 3) studying direct comparisons between human swarms and artificial swarms, as well as interaction between human swarm agents and artificial swarm agents, such as robots.

Acknowledgements. This work was partially supported by the program of Concerted Research Actions (ARC) of the Université libre de Bruxelles. M.K. Heinrich, A. Cleeremans and M. Dorigo acknowledge support from the F.R.S.-FNRS, of which they are, respectively, postdoctoral researcher and research directors.

References

1. Bailenson, J.N., Beall, A.C., Loomis, J., Blascovich, J., Turk, M.: Transformed social interaction: decoupling representation from behavior and form in collaborative virtual environments. Presence Teleoperators Virtual Environ. **13**(4), 428–441 (2004). https://doi.org/10.1162/1054746041944803
2. Barrington, L., et al.: Crowdsourcing earthquake damage assessment using remote sensing imagery. Ann. Geophys. **54**(6) (2011). https://doi.org/10.4401/ag-5324
3. Blascovich, J., Loomis, J., Beall, A., Swinth, K., Hoyt, C., Bailenson, J.: Immersive virtual environment technology as a methodological tool for social psychology. Psychol. Inq. **13**, 103–124 (2002)
4. Boos, M., Pritz, J., Lange, S., Belz, M.: Leadership in moving human groups. PLoS Comput. Biol. **10**(4), e1003541 (2014). https://doi.org/10.1371/journal.pcbi.1003541
5. Cheng, J.T., Tracy, J.L., Foulsham, T., Kingstone, A., Henrich, J.: Two ways to the top: evidence that dominance and prestige are distinct yet viable avenues to social rank and influence. J. Pers. Soc. Psychol. **104**(1), 103–125 (2013). https://doi.org/10.1037/a0030398
6. Cooper, S., et al.: Predicting protein structures with a multiplayer online game. Nature **466**(7307), 756–760 (2010). https://doi.org/10.1038/nature09304
7. De Montjoye, Y.A., Stopczynski, A., Shmueli, E., Pentland, A., Lehmann, S.: The strength of the strongest ties in collaborative problem solving. Sci. Rep. **4**, 5277 (2014)
8. Ducatelle, F., et al.: Cooperative navigation in robotic swarms. Swarm Intell. **8**(1), 1–33 (2013). https://doi.org/10.1007/s11721-013-0089-4
9. Eberhart, R., Palmer, D., Kirschenbaum, M.: Beyond computational intelligence: blended intelligence. In: 2015 Swarm/Human Blended Intelligence Workshop (SHBI). IEEE (2015). https://doi.org/10.1109/shbi.2015.7321679
10. Henrich, J., Heine, S.J., Norenzayan, A.: The weirdest people in the world? Behav. Brain Sci. **33**(2–3), 61–83 (2010). https://doi.org/10.1017/s0140525x0999152x
11. Heyes, C.: Who knows? Metacognitive social learning strategies. Trends Cogn. Sci. **20**(3), 204–213 (2016). https://doi.org/10.1016/j.tics.2015.12.007
12. Hunt, E.R., Jones, S., Hauert, S.: Testing the limits of pheromone stigmergy in high-density robot swarms. Roy. Soc. Open Sci. **6**(11), 190225 (2019). https://doi.org/10.1098/rsos.190225
13. Ioannou, C.C.: Swarm intelligence in fish? The difficulty in demonstrating distributed and self-organised collective intelligence in (some) animal groups. Behav. Process. **141**(2), 141–151 (2017)
14. Juliani, A., et al.: Unity: a general platform for intelligent agents. arXiv preprint arXiv:1809.02627 (2018). https://arxiv.org/pdf/1809.02627.pdf
15. Jung, J.H., Schneider, C., Valacich, J.: Enhancing the motivational affordance of information systems: the effects of real-time performance feedback and goal setting in group collaboration environments. Manage. Sci. **56**(4), 724–742 (2010). https://doi.org/10.1287/mnsc.1090.1129
16. Kalma, A.P., Visser, L., Peeters, A.: Sociable and aggressive dominance: personality differences in leadership style? Leadersh. Quart. **4**(1), 45–64 (1993). https://doi.org/10.1016/1048-9843(93)90003-c
17. Kendal, R.L., Boogert, N.J., Rendell, L., Laland, K.N., Webster, M., Jones, P.L.: Social learning strategies: bridge-building between fields. Trends Cogn. Sci. **22**(7), 651–665 (2018). https://doi.org/10.1016/j.tics.2018.04.003

18. Kirschenbaum, M., Palmer, D.W.: Perceptualization of particle swarm optimization. In: 2015 Swarm/Human Blended Intelligence Workshop (SHBI). IEEE (2015). https://doi.org/10.1109/shbi.2015.7321681
19. Krafft, P.M., et al.: Human collective intelligence as distributed Bayesian inference. arXiv preprint arXiv:1608.01987 (2016). https://arxiv.org/pdf/1608.01987.pdf
20. Krause, J., Ruxton, G.D., Krause, S.: Swarm intelligence in animals and humans. Trends in Ecol. Evol. **25**(1), 28–34 (2010). https://doi.org/10.1016/j.tree.2009.06.016
21. Kurvers, R.H.J.M., Wolf, M., Naguib, M., Krause, J.: Self-organized flexible leadership promotes collective intelligence in human groups. Roy. Soc. Open Sci. **2**(12), 150222 (2015). https://doi.org/10.1098/rsos.150222
22. Lepri, B., Staiano, J., Shmueli, E., Pianesi, F., Pentland, A.: The role of personality in shaping social networks and mediating behavioral change. User Model. User-Adap. Interact. **26**(2–3), 143–175 (2016). https://doi.org/10.1007/s11257-016-9173-y
23. Lin, A.Y.M., Huynh, A., Lanckriet, G., Barrington, L.: Crowdsourcing the unknown: the satellite search for Genghis Khan. PLoS ONE **9**(12), e114046 (2014). https://doi.org/10.1371/journal.pone.0114046
24. Mathews, N., Christensen, A.L., O'Grady, R., Mondada, F., Dorigo, M.: Mergeable nervous systems for robots. Nat. Commun. **8**(439) (2017). https://doi.org/10.1038/s41467-017-00109-2
25. Mekler, E.D., Brühlmann, F., Tuch, A.N., Opwis, K.: Towards understanding the effects of individual gamification elements on intrinsic motivation and performance. Comput. Hum. Behav. **71**, 525–534 (2017). https://doi.org/10.1016/j.chb.2015.08.048
26. Metcalf, L., Askay, D.A., Rosenberg, L.B.: Keeping humans in the loop: pooling knowledge through artificial swarm intelligence to improve business decision making. Calif. Manage. Rev. **61**(4), 84–109 (2019)
27. Michel, O.: Cyberbotics Ltd., Webots$^{\mathrm{TM}}$: professional mobile robot simulation. Int. J. Adv. Robot. Syst. **1**(1), 40–43 (2004)
28. Millard, A.G., et al.: The Pi-puck extension board: a Raspberry Pi interface for the e-puck robot platform. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 741–748. IEEE (2017)
29. Mondada, F., et al.: The e-puck, a robot designed for education in engineering. In: Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions, vol. 1, pp. 59–65. IPCB: Instituto Politécnico de Castelo Branco (2009)
30. Montes de Oca, M.A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., Dorigo, M.: Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. Swarm Intell. **5**(3–4), 305–327 (2011). https://doi.org/10.1007/s11721-011-0062-z
31. Moussaïd, M., Helbing, D., Garnier, S., Johansson, A., Combe, M., Theraulaz, G.: Experimental study of the behavioural mechanisms underlying self-organization in human crowds. Proc. Roy. Soc. B Biol. Sci. **276**(1668), 2755–2762 (2009)
32. Moussaïd, M., et al.: Crowd behaviour during high-stress evacuations in an immersive virtual environment. J. Roy. Soc. Interface **13**(122), 20160414 (2016). https://doi.org/10.1098/rsif.2016.0414
33. Mulders, D., De Bodt, C., Bjelland, J., Pentland, A., Verleysen, M., de Montjoye, Y.A.: Inference of node attributes from social network assortativity. Neural Comput. Appl. 1–21 (2019). https://doi.org/10.1007/s00521-018-03967-z

34. Nakamura, J., Csikszentmihalyi, M.: The concept of flow. Flow and the Foundations of Positive Psychology, pp. 239–263. Springer, Dordrecht (2014). https://doi.org/10.1007/978-94-017-9088-8_16
35. Pedersen, M.K., Rasmussen, N.R., Sherson, J.F., Basaiawmoit, R.V.: Leaderboard effects on player performance in a citizen science game. In: Proceedings of the 11th European Conference on Game Based Learning, vol. 531 (2017)
36. Pinciroli, C., Talamali, M.S., Reina, A., Marshall, J.A.R., Trianni, V.: Simulating Kilobots within ARGoS: models and experimental validation. In: Dorigo, M., Birattari, M., Blum, C., Christensen, A.L., Reina, A., Trianni, V. (eds.) ANTS 2018. LNCS, vol. 11172, pp. 176–187. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00533-7_14
37. Pinciroli, C., et al.: ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. Swarm Intell. **6**(4), 271–295 (2012). https://doi.org/10.1007/s11721-012-0072-5
38. Quigley, M., et al.: ROS: an open-source robot operating system. In: ICRA Workshop on Open Source Software, vol. 3, p. 5, Kobe, Japan (2009)
39. Quinn, A.J., Bederson, B.B.: Human computation: a survey and taxonomy of a growing field. In: Proceedings of the International Conference on Human Factors in Computing Systems (2011)
40. Rosenberg, L., Baltaxe, D., Pescetelli, N.: Crowds vs swarms, a comparison of intelligence. In: 2016 Swarm/Human Blended Intelligence Workshop (SHBI). IEEE (2016). https://doi.org/10.1109/shbi.2016.7780278
41. Rosenberg, L.B.: Human swarms, a real-time method for collective intelligence. In: 20/07/2015–24/07/2015. The MIT Press (2015). https://doi.org/10.7551/978-0-262-33027-5-ch117
42. Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. Science **345**(6198), 795–799 (2014). https://doi.org/10.1126/science.1254295
43. Sørensen, J.J.W.H., et al.: Exploring the quantum speed limit with computer games. Nature **532**(7598), 210–213 (2016). https://doi.org/10.1038/nature17620
44. Sørensen, J.J.W., et al.: Exploring the quantum speed limit with computer games. Nature **532**(7598), 210–213 (2016)
45. Thrash, T., et al.: Evaluation of control interfaces for desktop virtual environments. Presence Teleoperators Virtual Environ. **24**(4), 322–334 (2015). https://doi.org/10.1162/pres_a_00237
46. Valentini, G., Ferrante, E., Dorigo, M.: The best-of-n problem in robot swarms: formalization, state of the art, and novel perspectives. Front. Robot. AI **4** (2017). https://doi.org/10.3389/frobt.2017.00009
47. Valentini, G., Hamann, H., Dorigo, M.: Self-organized collective decision-making in a 100-robot swarm. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI 2015), pp. 4216–4217. AAAI Press (2015)
48. Vasile, C., Pavel, A., Buiu, C.: Integrating human swarm interaction in a distributed robotic control system. In: 2011 IEEE International Conference on Automation Science and Engineering, pp. 743–748. IEEE (2011)
49. Walker, P., Amraii, S.A., Chakraborty, N., Lewis, M., Sycara, K.: Human control of robot swarms with dynamic leaders. In: 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1108–1113. IEEE (2014)
50. Wang, H., Sun, C.T.: Game reward systems: gaming experiences and social meanings. In: Proceedings of DiGRA 2011 Conference: Think Design Play (2012)

51. Werfel, J., Petersen, K., Nagpal, R.: Designing collective behavior in a termite-inspired robot construction team. Science **343**(6172), 754–758 (2014). https://doi.org/10.1126/science.1245842
52. Whiten, A., Hinde, R.A., Laland, K.N., Stringer, C.B.: Culture evolves. Philos. Trans. Roy. Soc. B Biol. Sci. **366**(1567), 938–948 (2011). https://doi.org/10.1098/rstb.2010.0372
53. Zhao, H., et al.: A networked desktop virtual reality setup for decision science and navigation experiments with multiple participants. J. Vis. Exp. **138**(e58155) (2018). https://doi.org/10.3791/58155