Authors' accepted manuscript. Article to be published in *IEEE Robotics & Automation Magazine*.

Planetary exploration with robot teams

David St-Onge[1], Marcel Kaufmann[1], Jacopo Panerati[1], Benjamin Ramtoula[1], Yanjun Cao[1], Emily B.J. Coffey[2] and Giovanni Beltrame[1]
[1] Department of Computer Engineering and Software Engineering, Polytechnique Montreal, Canada
[2]Department of Psychology, Concordia University, Canada

Since the beginning of space exploration, Mars and the Moon have been explored with orbiters, landers, and rovers. Over forty missions have targeted Mars, and more than a hundred, the Moon. Developing novel strategies and technologies for exploring celestial bodies continues to be a focus of space agencies. Multi-robot systems are particularly promising for planetary exploration, as they are more robust to individual failure and have the potential to explore larger areas; however, there are limits to how many robots an operator can individually control. We recently took part in the European Space Agency's interdisciplinary equipment test campaign (PANGAEA-X) at a Lunar/Mars analogue site in Lanzarote, Spain. We used a heterogeneous fleet of Unmanned Aerial Vehicles (UAVs)—a swarm—to study the interplay of systems operations and human factors. Human operators directed the swarm via ad-hoc networks and data sharing protocols to explore unknown areas under two control modes: one in which the operator instructed each robot separately; and the other in which the operator provided general guidance to the swarm, which self-organized via a combination of distributed decision-making, and consensus building. We assessed cognitive load via pupillometry for each condition, and perceived task demand and intuitiveness via self-report. Our results show that implementing higher autonomy with swarm intelligence can reduce workload, freeing the operator for other tasks such as overseeing strategy, and communication. Future work will further leverage advances in swarm intelligence for exploration missions.

***Index Terms*—space exploration, decentralized robotics, unmanned aerial vehicles, human-swarm interaction**

I. INTRODUCTION

Extraterrestrial exploration missions are increasingly directed towards challenging landscapes and environments, such as mountains, craters, lava tubes, and oceans [1], [2], [3], [4]. Rovers have been the most common vehicle choice for planetary exploration to date. However, they are designed to operate on relatively flat land [5]. As a result, other types of robots such as unmanned aerial vehicles (UAVs) and hydrobots are being proposed for more challenging environments and geographic features, including the atmospheres and oceans of celestial bodies [6]. UAVs offer advantages over rovers for exploration where atmospheres are present: they provide higher resolution data as compared to orbiters [7], greater range and mobility [6], and they can sample gases at different altitudes, thus also filling a planetary measurement gap [8].

Multi-robot teams, perhaps with mixed capacities (i.e. activators, sensors and communication devices), could also be used to explore larger areas more effectively than single robots, and could characterize and identify potential landing sites for manned missions, as well as reveal hazardous areas. With an appropriate interface, a robotic team could conduct autonomous reconnaissance [9] and increase human situational awareness of mission-critical information. However, there are substantial technical and human challenges to organizing and controlling multi-robot systems.

The European Space Agency recently invited our team (Fig. 1) to run an experiment on the use of a multi-robot aerial system for planetary exploration as part of their PANGAEA-X exercise, a test campaign that brings together astronauts, scientists, engineers and operations experts for advancing integrated human and robotics missions. Participants, including a European astronaut, controlled a heterogeneous fleet of four to six UAVs.

Our first objective was to demonstrate the physical deployment of the UAV team. Long distances generate communication latencies and impose low bandwidth, so we rely on decentralized control for our robots. As such, all UAVs are replaceable by any others, improving the robustness of the overall system to individual robot failures. Our approach is suitable for gathering aerial images and to provide the operators with a fleet-wide communication link over kilometres, under challenging real-world conditions. These characteristics are made possible by the combination of several of our core contributions to swarm robotics and to a novel approach of Human-Swarm Interaction: our system sees the operator as just another robot.

The PANGAEA-X context presented a rare occasion to measure human behaviour in an operational environment. Task performance and risk-taking behaviour [10] differ in the field as compared with simulated robotic tasks, likely because real situations are more engaging and potentially stressful. Our second objective was therefore to study the human operator as they guide the swarm. We addressed two questions: (i) How does the operator's perception of usability and workload change over levels of swarm autonomy? (ii) How do objective measures of user cognitive workload change over levels of autonomy?

To answer these questions, we created two control modes for our swarm that differ by degree of robotic autonomy. We monitored the operator's cognitive load via pupillometry during task performance, and subsequently assessed the user's subjective experience. This paper presents the main components and contributions of our experimental field setup and discusses the human-swarm interaction results obtained at PANGAEA-X.

Manuscript received XX; revised XX, 2019. Corresponding author: D. St-Onge (email: david.st-onge@polymtl.ca).

Fig. 1. PANGAEA-X field deployment team in Lanzarote, Spain: five engineers, a neuroscientist, five DJI Matrice 100 and five Pleiades Spiri. The volcanic landscape of Lanzarote is similar to the surface of the Moon. (Image: ESA)

II. Decentralized control

In safety and mission-critical multi-robot applications, ending a mission because of a single unit failure is unacceptable. For example, when a group of rovers explores a lava tube on the Moon, if the leader robot—the one coordinating the mission and allocating the tasks—gets stuck, the team's exploration potential will be greatly reduced. In an emergency response scenario in which a team of aircraft are searching for victims, losing the link to the control station responsible for trajectory planning could result in the loss of life. Decentralized paradigms, which use only local information for control and communication, can mitigate these issues while rendering a team more adaptable to dynamic environments. For instance, failure of a single robot or a broken communication link would not compromise the mission as the remaining robots can collaboratively reorganize their activities to cover the search region.

For these reasons, we consider decentralized control to be an ideal solution for space applications, as previously shown for the formation control of multiple collaborating spacecraft [11], or to synchronize the actuators of a Martian ground robot [12].

Figure 2 shows the details of our implementation for these experiments. For inter-robot communication, we rely on Xbee mesh modules and for localization, on GPS. Because the robots' coordination relies solely on inter-robot distance and bearing, localization is required only for user inputs. GPS can be substituted with other available means (e.g. ultra-wide band, landmarks, or camera-based mapping).

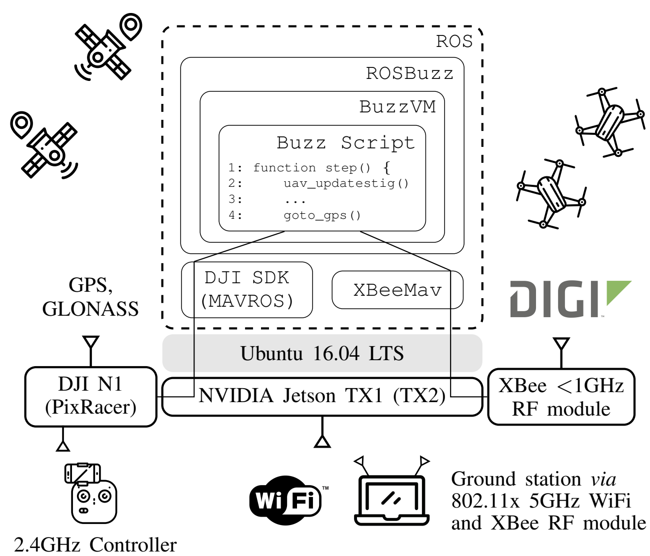

Fig. 2. The control architecture (from left to right): the localization system and (back-up) remote controllers interface with the UAVs' onboard computers and flight controllers (FCU) to execute the decentralized behavioural Buzz script. The entire fleet executes the same script, interfacing with the FCU and communication device (Xbee) through the Robot Operating System (ROS) and Mavlink protocol.

A. Swarm behaviour design

Developing sophisticated, fully decentralized behaviours is very challenging, as they can only rely on limited information and local interactions. To simplify our implementation, we use Buzz, a domain-specific programming language for robot swarms [13]. Buzz provides special constructs for robots to share data with the swarm (a technique called *virtual stigmergy*) and interact with their neighbours. Buzz is highly portable as it runs within a minimalist virtual machine that works on most computing systems. Buzz merges bottom-up behaviour development (i.e. assigning tasks to specific robots) with top-down programs controlling the whole swarm. The developer using Buzz can implement high-level coordination algorithms while still considering the specificity of each of the units in a heterogeneous team, such that the same code can be deployed on any autonomous robot. The algorithms implemented in Buzz for this experiment were tested in a Robot Operating System (ROS) environment with the Gazebo simulator, before field deployment. The Buzz virtual machine [14], as well as its ROS integration illustrated in Fig. 2—including the script files used in this work—are freely available online [15].

B. Communication with neighbours

In critical application scenarios, it can be complicated to maintain a reliable connection to all the robots in a team. Multiple challenges compound: (i) large areas to cover, (ii) limited 1-hop communication ranges and (iii) the (human) command centre being potentially located in a remote location for safety or operational reasons. While a central control architecture requires a link to each robot, decentralized paradigms can support many other sparse network topologies. Decentralized control only requires one of the robots to be within communication range to the ground station to send commands and receive status updates (using a Wi-Fi mesh, Xbee, Zigbee, etc.). The information can then be propagated, gossip-like, from robot to robot (as shown in Fig. 3*D*). Each robot's memory contains two tables. The "neighbours" table includes the virtual machine state, as well as sensor output such as GPS (confidence and coordinates), network device state (link quality), and battery level. The "mission" table includes mission-relevant information such as the next goal(s) or the location of landmarks of interest. Both tables are reliably shared across the robot team through Buzz's virtual stigmergy. In a nutshell, when a piece of data is required by an individual robot, the team agrees on who has its latest version and the requester receives the updated value through the shortest network path available (i.e. the one with fewest hops). This strategy optimizes the use of the bandwidth, allowing the fleet to rely on long-range transmission devices with low data rates. In previous experiments [16], we stressed our communication architecture and demonstrated that our consensus strategy and information-sharing protocol were resilient with up to 80% packet loss. Field tests showed that we can rely on inter-robot, 1-hop links of over 600 metres using 800-900 MHz radios (Digimesh Xbee S3B Pro modules). The relations between communication device, Buzz virtual machine, ROS and the localization system are illustrated in Fig. 2.

C. Safety features

In a remote operation scenario for planetary exploration sudden changes in the environment may arise too quickly for the operator to send each robot a timely command. To cope with bursts of wind, for instance, we use a virtual GPS fence–the geofence–according to limits set before take-off (Fig. 3*G*). Additional layers of security are required to cope with communication and human errors. For example, when setting a goal, the destination's distance is verified to be reachable in less than 30 seconds (e.g., less than 300 metres, if the UAV maximum velocity is 10m/s). This verification is done on input, before sending the command, and again by the receiving robot. As the robots generate their own trajectories, they can use the known locations of their neighbours to avoid each other, using a decentralized collision avoidance algorithm [17]. To minimize risk of collision, we also flew UAVs at different altitudes.

For each robot, the potential sources of failure increase with its mechanical, communication and instrumentation complexity. If a robot self-detects an imminent failure, it can upload its current state and mission role to the fleet's shared memory using a mechanism derived from the virtual stigmergy for larger data [18]. Its identity will then be accessible for another available unit to take over its mission. To achieve system-wide robustness, we wrapped a consensus strategy around the various critical steps of the mission. Before take-off and before accepting user inputs, all members of the team verify that enough units are available and ready to execute the next task. If a UAV does not get acknowledgment from its expected number of neighbours, it hovers, waiting up to two minutes for the missing teammate(s) to show up. If the issue is not resolved, the hovering drone will broadcast a message to let the team know that the next step cannot be processed in the current state. Thereupon, the robot will either stay in its previous state or proceed to the landing zone, according to the operator's preference. In that regard, the operator is also considered a member of the fleet, and if the link is broken, the same procedure will be triggered. The safety mechanisms discussed in this section are illustrated in Fig. 3(*G*).

III. Operator command and control

A command centre for exploration missions must ease the optimization of strategic allocation of field resources, situational awareness and collaboration between team members. For deployed robotic systems, the command centre design is adapted to the robots' level of autonomy. Most space missions fall between teleoperation and fully autonomous missions, having several scripted behaviours as well as high-level commands [19]. Specific to robot teams, novel interaction modalities have been studied for shared autonomy: using gesture control [20], voice [21], and even full-body motion [22]; however, maplike interfaces are popular for critical scenarios as they are familiar and leverage operators' existing skills [23]. For UAVs, a map-based interface is commonly referred to as a *mission planner*: an application running on a ground station monitoring the fleet. Currently available commercial mission planners (DroneDeploy, DJI Go, QGroundControl, etc.) integrate maps and point-and-click waypoint selection to ease route design for operators. Fig. 4-a is a screenshot of the QGroundControl mission planning software. As shown in Fig. 4-b, our interface is also based on a map service integrated in a browser control panel. This interface allows for

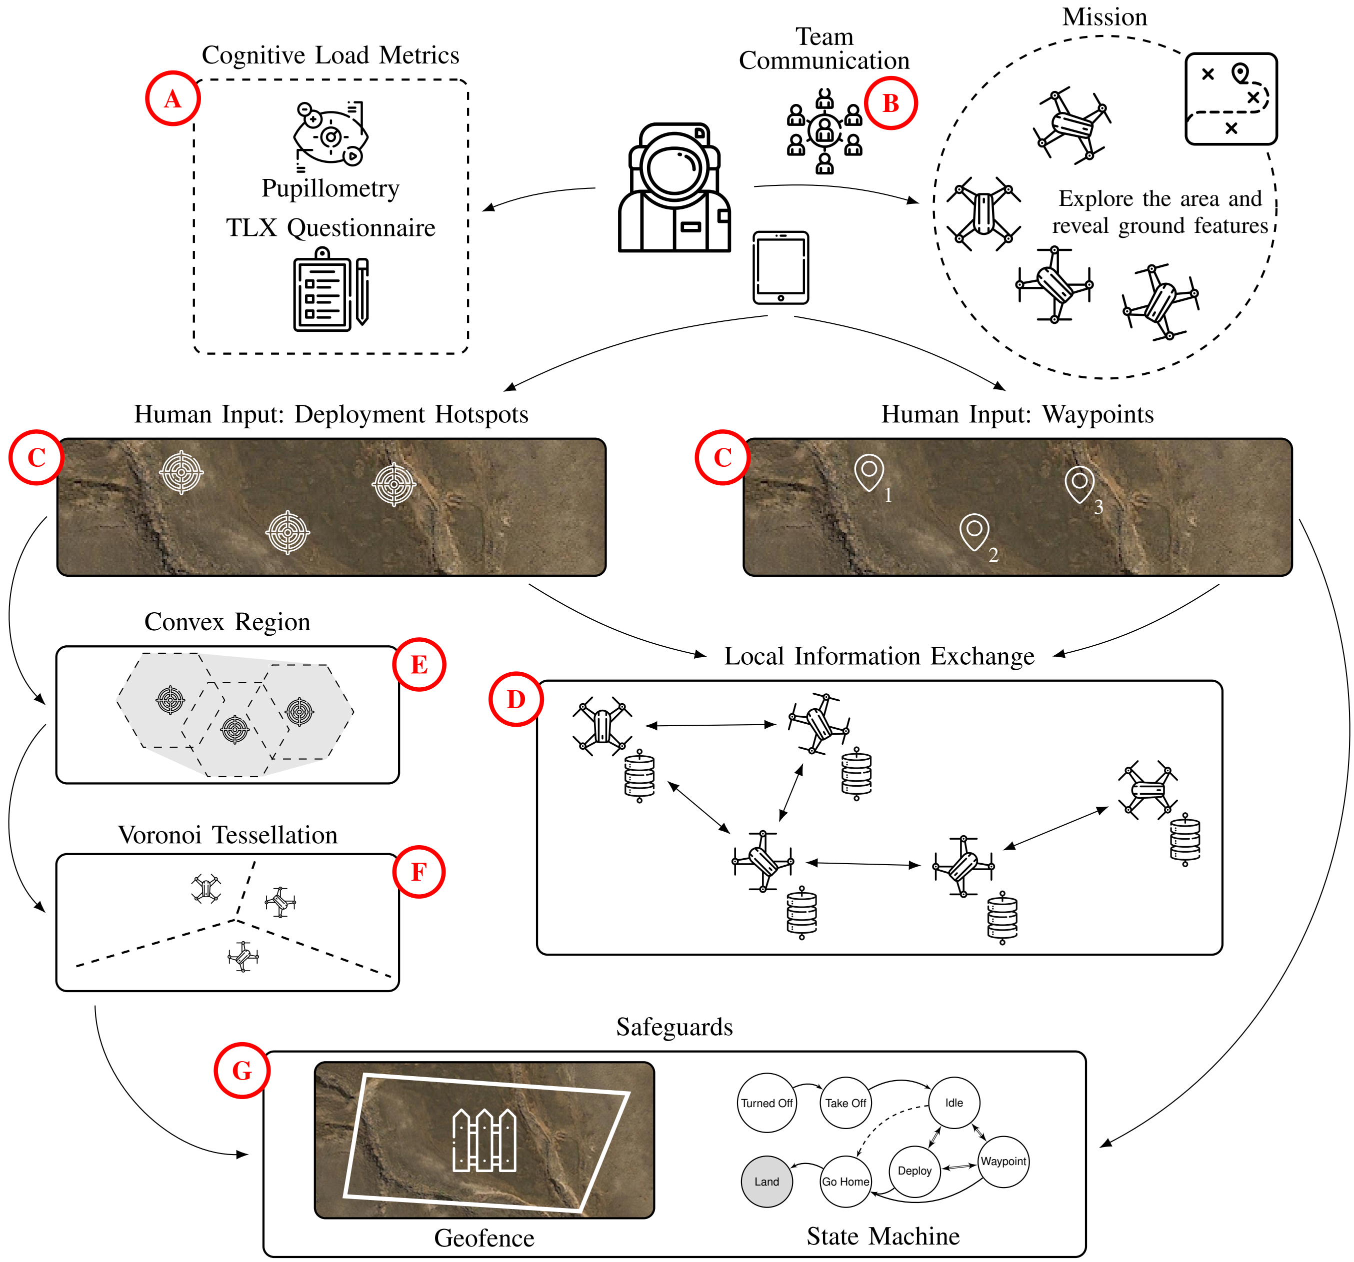

Fig. 3. Experiment overview: A) Objective and self-assessed metrics of the operator cognitive load, B) Simulated radio chatter to increase operator task load, C) Operator inputs following the two modes (hotspots for self-deployment and waypoints for individual control), D) Decentralized communication infrastructure, E) Convex region computation on each robot from the user hotspots input, F) Location goals computed from the tesselation of the region of interest and G) Safety features implemented (geofence set before the mission and a state machine with consensus mechanism over transitions).

the operator to monitor detailed aspects of each robot (battery level, current state and position), send general commands (e.g. take off or go home) and specific commands related to the fleet control mode: waypoint selection or fleet deployment. In all scenarios, the command centre only acts as a member of the swarm showing the other members' state and sending updates to the virtual stigmergy data; there is no centralized control. Path planning and collision avoidance are computed on-board each UAV.

A. *Waypoint Selection Mode*

When operating under the "waypoint selection mode", our command centre displays many of the features that are common to other mission planners, such as the one in Fig. 4. In this mode, each robot is controlled individually: the operator clicks on the map and selects the desired robot from a pop-up menu, generating a target waypoint. This is repeated for every robot, resulting in individual goals (Fig. 3*C*). The waypoint command generated is sent to the closest UAV from the ground station, and propagated throughout the fleet. While the mission planning strategy is centralized around the user sending com-

(a) (b)

Fig. 4. Mission planers screenshot: a) A popular interface, QGroundControl, with five waypoints defined for an UAV route, b) our web-based minimal interface with a waypoint popup menu (individual waypoint control).

mands, communication is decentralized to ensure tolerance to link and robot failures [24]. With our infrastructure, each waypoint can also be attributed to any other UAV, since the list is shared among the fleet. This falls into a task allocation optimisation problem, with several solution in the literature, for instance [25], which ensures that a goal not reached by its associated robot will be back in an active list of goals for another one to pick.

At any time during the mission, the user may clear the waypoint list, triggering the fleet to switch to its hovering state until further instructions arrive. The user may also dynamically change the waypoints, even if the targeted UAV has not yet reached its goal. By doing so, a fast operator can act as if he or she is tele-operating all the UAVs simultaneously. However, dynamic individual control requires constant monitoring and input. The cognitive resources of the operator thus become a limitation on swarm size, as each additional robot receives less individual attention and may be neglected and become ineffective for part of its exploration time [26]. Furthermore, additional operator tasks, such as communicating with (human) teammates, increases the operator's cognitive load [27], leading to slower deployment of robots.

B. *Self-deployment mode*

In the more autonomous mode, the operator identifies several "hotspots" that may be worth exploring. This control mode exploits concepts from computational geometry to implement enhanced autonomy within the fleet. The locations received from the operator are shared across the whole fleet and each robot computes a minimum convex polygonal region of interest (Fig. 3*E*). This region is then split into *cells* to be distributed among the swarm members, a process known as surface tessellation. Some applications such as search and rescue and sensor network deployment already use similar approaches [28]. For a uniform coverage of the area to be explored, we select the Voronoi tessellation, for which cells contain all points that are geometrically closer to their centre. For deploying robotic teams, the initial positions can be taken as *seeds* to the tessellation problem. A possible implementation consists of creating a frontier (line) halfway between each 'seed' (e.g. robot) and merging these lines into polygon edges. Distributed computation of the Voronoi tessellation was extensively studied for multi-robot deployment [29]. We use the *sweeping line algorithm* (Fortune's algorithm), which is one of the most efficient ways to extract cell lines from a set of seeds [30]. We then cut the open cells using the user-defined convex polygonal boundary. From this point on, each robot has knowledge of its Voronoi cell's limits. To reach a uniform distribution of the robots in the area, we use a simple gradient descent towards the centroid of each cell [28]. If one of the robots is not yet in the region of interest, the robot takes a random location goal inside the zone and becomes a seed for the uniform deployment as soon as it enters the region. Each robot recomputes the tessellation following updates on the relative position of its neighbours; an approach that is robust to both packet loss and dynamic inputs. The operator can change the shape and location of the region of interest at any time, and the robots adapt accordingly.

IV. Operator cognitive load

Controlling larger numbers of robots is likely to increase the cognitive load of the operator [31]. Cumulative cognitive load might be a function of attention required to supervise the robots [32], a robot's performance and autonomy when left unattended [33], and the need to manage dependencies between robots when they must coordinate to perform a task [34], in addition to other mission requirements such as communication with teammates. The field of Human-Swarm Interaction (HSI) specifically addresses the tension between a central element of control and a decentralized system [35]: a human operator issues commands to a swarm that may dynamically organize its configuration and the interdependencies between its robots. In literature, most HSI studies are conducted in simulation. However, human behaviour in real-world applications differs as compared to in simulation [36], suggesting that the operator's cognitive activities might also differ contextually. To avoid the issue, we conduct experiments with physical robots.

Fig. 5. German astronaut Matthias Maurer wearing eye-tracking glasses (Pupil Labs) while conducting an exploration mission. He holds a tablet running the mission planner and observes the UAVs' reaction to input. In background, team members offer suggestions during familiarization training, and stand by to assume manual control of the UAVs in case of an emergency. (Image: ESA)

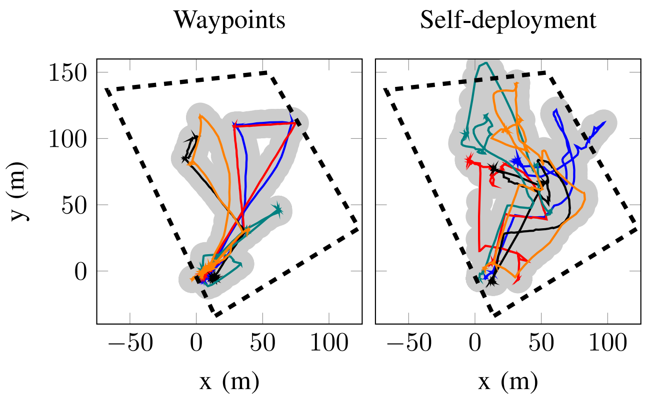

Fig. 6. Top views of five UAV trajectories (different colours) comparing the two control modes from a representative participant. The more autonomous self-deployment condition (right) shows better area coverage. The more manual waypoint condition (left) shows duplicated paths for different UAVs. UAVs were twice pushed beyond the geofence by strong winds in the self-deployment condition (right).

Our experiment would not be realistic if we ignored all the other tasks that human operators would need to perform in parallel for planetary exploration. We created fictional team radio chatter, which is played over earphones (Fig. 3*B*). The audio files contain contextual information such as "John is going out on EVA, keep an eye out", and occasional mission-specific questions prefaced with the call sign "Operator," such as "how far is Robot One from the initial point?" The operator is instructed to quickly acknowledge the communication with a button press on the control screen, and respond verbally. In this initial sample, the supplementary task serves to ensure that the speech is attended, thus increasing workload; in larger sample sizes it would be possible to use missed vs. correct responses as an additional performance measure for statistically comparing conditions.

Cognitive load can be measured through subjective self-assessment metrics such as questionnaires; and objective metrics such as body motion, heart rate variability, and measures of pupil dilation over time derived from pupillometry [37], [38]. Pupillometry has recently gained popularity in applied psychology as a reliable proxy for cognitive load [39]. For example, pupil dilation was used to study the effects of audiovisual interference on workload in piloting tasks [40].

We use both types of measurements: (i) a questionnaire, which included questions inspired by a survey originally designed to evaluate the perceived usability and acceptance of assistive devices [41], augmented with task-load oriented questions from the NASA Task-Load Index (TLX); and (ii) pupil dilation and variability (Fig. 3*A*). To provide equivalent psychological distance between the scores [42], all answers were on a seven-point Likert scale (0 to 6).

V. RESULTS: PANGAEA-X FIELD TEST

Our field experiment was conducted in Lanzarote, Spain. The unique landscape of this volcanic island is one of the closest to a lunar landscape one can find on Earth, and is used for geology training for astronauts by ESA.

Our sample is small (five participants) due to logistic and weather constraints: setting up the UAV fleet took ∼30 minutes per run, each trial lasted about an hour, and UAVs have weather mimina to safety operate (visibility, light, precipitation, and wind speed). We first performed a series of experiments on our own team members (beta testers), then we conducted a core set of experiments on two members of the ESA crew (men) and one journalist (woman).

Informed consent was obtained. The eye tracker was calibrated [43], then the operators explored an area for hidden ground features by guiding the UAVs using each of the two control modes, in separate missions, while listening and responding to intermittent radio chatter. After each of the two missions, the participants' feedback on their user experience was collected.

Exploration task performance

The aim of each mission was to search for hidden ground features - a task best accomplished by widely covering the search area. An representative example of the UAVs' trajectories for the two control modes is shown in Fig. 6. The area covered in self-deployment mode is larger than in the waypoint mode, whereas in the waypoint mode there are duplicate trajectories, indicating lower overall search efficiency. This outcome is expected since the self-deployment algorithm aims at spreading the robots over the area of operator-defined interest. In the waypoint mode, the human operators must place the waypoints to sweep the area, while mentally keeping track of the area that each robot has already explored.

On average, for both control modes the participants discovered most (three out of the five total) of the hidden ground features; testing a larger sample may reveal differences. Anecdotally, most participants expressed confusion about which areas had already been explored in the waypoint condition, and they also relied on their memory to assign new goals to the UAVs (instead of clicking on a UAV to get its ID), often sending an unintended UAV towards a new goal. Similar observations apply to the trajectories obtained by the other participants.

Perception of usability and workload over levels of swarm autonomy

Figure 7 presents the average results of the main survey elements for both control modes. For both, the interface was considered similarly easy to learn (> 4.6), easy to use (> 4.4), intuitive (> 4.4) and effective ($= 4.4$) without being cumbersome (< 1.5). We attribute these promising initial results to the minimalist and clutter-free design of our mission planner (see Fig. 4-b) and to the good will of the participants. The results also show that neither task was considered highly demanding (< 3.0) or hard to complete (< 3.0).

Figure 7 shows that individual waypoint control gives the operator more confidence (less insecurity) about the outcome of his or her actions. Results also show that the self-deployment mode was less intuitive, less efficient, and harder to learn. We can also observe that this control mode was slightly more mentally demanding. We believe these per-

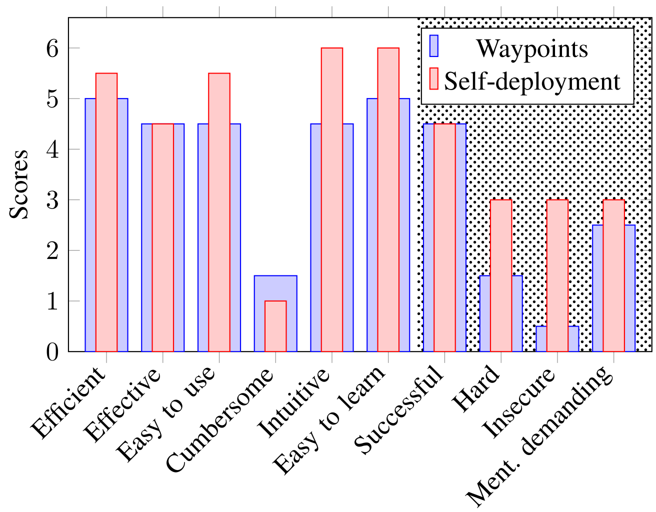

Fig. 7. Average results from the questionnaire comparing the perceived usability (left) and perceived task load (right) of each control mode.

ceptions are related to the impression of not being fully in control; in the self-deployment mode, the relations between the operator's intentions and the robots' behaviour is not explicit. Such self-organized and emerging behaviours are more challenging to visualize than deterministic control strategies. These observations suggest that further developments in mission planner data representation will be needed to best communicate from swarm to operator. Furthermore, results suggest that the users' impression of control may influence their perception of cognitive load.

Objective measures of cognitive workload over levels of swarm autonomy

An external objective measure is required to assess the real difference in cognitive load between the control modes. Pupillometry data was processed to remove outliers ($> \frac{\sigma}{2}$), subtract the median diameter, and was low-pass filtered. The resulting curves of the two missions for the same participant are shown in Fig. 8. We overlay on the image the periods with background radio chatter (orange), direct verbal questions to the operator (green) and operator responses (red).

The pupillometry metrics used here in fact lead to different conclusions than the user's self-assessment. The results show that the operator's pupil dilation range is on average higher for the waypoint control, suggesting a higher cognitive load for this control mode. The background radio chatter has similar effect on both missions. Interestingly, both modes led to half of the acknowledgments of direct questions being missed, most likely because the operators were already overwhelmed with the mission.

VI. Conclusion

The ESA PANGAEA-X field campaign provided our team with a unique opportunity to test a swarm fleet in a realistic planetary analogue scenario. Leveraging our expertise on decentralized behaviour design, simulation-to-field software workflow and resilient robotic teams, we deployed a heterogeneous fleet of UAVs to demonstrate the robustness of the setup and the study protocol, which will support future development; and to compare the operator's experience of two different levels of embedded autonomy. We demonstrated that the technology of fully decentralized robotic systems is mature enough to provide a robust basis for human factor studies in the field - a first for robotics swarms. The small sample of operators involved does not guarantee the generalization of the human findings, but does suggest that both objective and subjective measures will help us to understand and improve human-swarm interactions. From the subjective measures, we learned that managing up to five UAVs simultaneously is still feasible for both the more and less autonomous control modes; greater benefits of autonomy are likely to arise as systems are scaled up. We also observed that the uncertainty of the output from the swarm's self-organized behavior led to operator confusion and insecurity, and greater perceived cognitive load. It appears that workload can partly be disentangled from perceived control using objective physiological measures, such as pupillometry. The pupil dilation data seem to indicate that the mode with greater autonomy (self-deployment) is indeed less demanding. Nonetheless, the results point to a need for interfaces that communicate prospective swarm behaviour more clearly if the full potential of autonomous swarms is to be realized. Our results motivate work on more intuitive command centres such as new graphic overlays of prospective swarm behaviour on the mission planner, or a tangible table interface replacing the mission planer entirely. We believe that swarm intelligence deployment in planetary exploration missions has a promising future, and that the technology will also be applicable to many Earth-based mobile autonomous systems.

We are currently conducting a larger experiment to allow for statistical comparisons between conditions. While still using swarms of real robots, a miniature indoor setup will require less preparation time and will ensure more stable environmental and luminance conditions, facilitating human measurements. We will also complement pupillometry with measures of skin conductivity and heart rate variability, which have different physiological bases and will provide a more complete perspective on the operator's state. To firmly anchor our contribution to the development of technologies for space exploration, we are now adapting our tools for ground-to-air robotic teams in lava tubes. These regions are difficult to observe from space but may well be the most suitable environments for safe settlements on Moon and Mars [44].

Acknowledgments

The authors would like to acknowledge the invitation of the European Space Agency and the support of the PANGAEA-X logistics team. Part of this research is funded by NSERC grants and half the fleet was designed and adapted to our needs by Pleiades Robotics.

References

[1] J. Bleacher, P. Richardson, W. Garry, J. Zimbelman, D. Williams, and T. Orr, "Identifying lava tubes and their products on olympus mons,

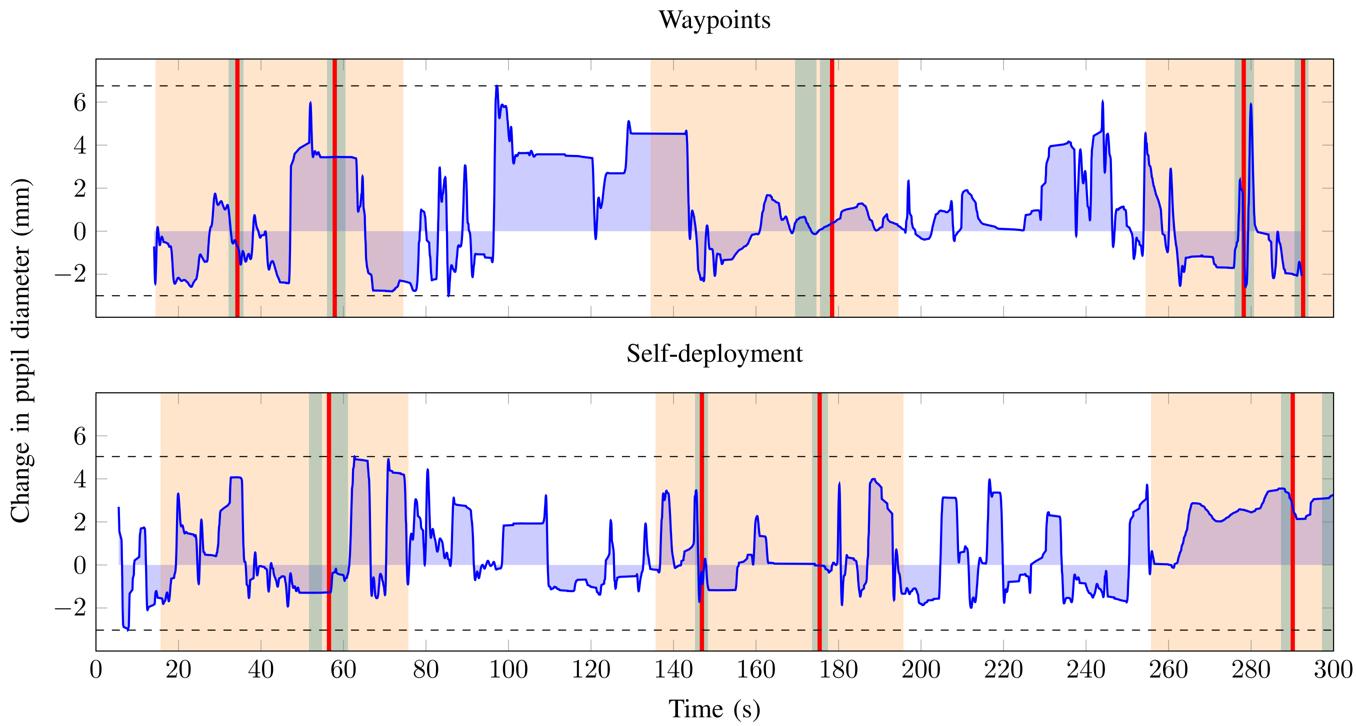

Fig. 8. Pupil diameter variations (from the median value) for both control modes for a representative participant. Periods in which the operator is required to attend to background discussion are shown in orange. Green bars indicate direct questions to the operator, and red bars indicate the operator's acknowledgement. Overall, the self-deployment mode shows less pupil diameter variation, suggesting a lighter cognitive load [38].

mars and implications for planetary exploration," in *Lunar and Planetary Science Conference*, vol. 42, 2011, p. 1805.

[2] M. Cardinale, S. Silvestro, D. A. Vaz, T. Michaels, M. C. Bourke, G. Komatsu, and L. Marinangeli, "Present-day aeolian activity in herschel crater, mars," *Icarus*, vol. 265, pp. 139–148, 2016.

[3] R. W. Zurek, J. R. Barnes, R. M. Haberle, J. B. Pollack, J. E. Tillman, and C. B. Leovy, "Dynamics of the atmosphere of mars," *Mars*, pp. 835–933, 1992.

[4] A. Fortes, "Exobiological implications of a possible ammonia–water ocean inside titan," *Icarus*, vol. 146, no. 2, pp. 444–452, 2000.

[5] T. D. J. M. Sanguino, "50 years of rovers for planetary exploration : A retrospective review for future directions," *Robotics and Autonomous Systems*, vol. 94, pp. 172–185, 2017.

[6] M. Hassanalian, D. Rice, and A. Abdelkefi, "Evolution of space drones for planetary exploration : A review," *Progress in Aerospace Sciences*, vol. 97, pp. 61–105, 2018.

[7] A. Elfes, S. S. Bueno, M. Bergerman, E. C. de Paiva, J. G. Ramos, and J. R. Azinheira, "Robotic airships for exploration of planetary bodies with an atmosphere: Autonomy challenges," *Autonomous Robots*, vol. 14, no. 2, pp. 147–164, Mar 2003.

[8] A. A. Gonzales, C. J. Cropus, D. W. Hall, and R. W. Parks, "DEVELOPMENT OF A USEFUL MARS AIRPLANE EXPLORATION CONCEPT AT NASA / AMES RESEARCH CENTER," in *6th Mars Society Conference*, 2003.

[9] NASA, "Mars airplane," Last visited 03/2019.

[10] J. W. Rudolph, D. B. Raemer, and R. Simon, "Establishing a Safe Container for Learning in Simulation," *Simulation in Healthcare: Journal of the Society for Simulation in Healthcare*, vol. 9, no. 6, pp. 339–349, 2014.

[11] M. Nazari and E. A. Butcher, "Decentralized Consensus Control of a Rigid-Body Spacecraft Formation with Communication Delay," *JOURNAL OF GUIDANCE,CONTROL, AND DYNAMICS*, vol. 39, no. 4, pp. 1–14, 2016.

[12] J. Bruce, "Design, Building, Testing, and Control of SUPERball: A Tensegrity Robot to Enable New Forms of Planetary Exploration," PhD thesis, University of California, 2016.

[13] C. Pinciroli and G. Beltrame, "Swarm-Oriented Programming of Distributed Robot Networks," *Computer*, vol. 49, no. 12, pp. 32–41, 2016.

[14] MIST Laboratory, "Main buzz virtual machine repository," Last visited 01/2019. [Online]. Available: https://github.com/MISTLab/Buzz

[15] ——, "Ros node to use buzz virtual machine," Last visited 01/2019. [Online]. Available: https://github.com/MISTLab/ROSBuzz

[16] D. St-Onge, V. S. Varadharajan, G. Li, I. Svogor, and G. Beltrame, "ROS and buzz: consensus-based behaviors for heterogeneous teams," *CoRR*, vol. abs/1710.08843, 2017. [Online]. Available: http://arxiv.org/abs/1710.08843

[17] M. Shahriari, I. Svogor, D. St-Onge, and G. Beltrame, "Lightweight collision avoidance for resource-constrained robots," in *IEEE/RSJ Conference on Intelligent Robots (IROS)*, 2018.

[18] V. S. Varadharajan, D. St-Onge, B. Adams, and G. Beltrame, "Soul: data sharing for robot swarms," *Autonomous Robots*, May 2019.

[19] V. Dimitrov and T. Padir, "A comparative study of teleoperated and autonomous task completion for sample return rover missions," *IEEE Aerospace Conference Proceedings*, 2014.

[20] E. A. Cappo, A. Desai, and N. Michael, "Robust Coordinated Aerial Deployments for Theatrical Applications Given Online User Interaction via Behavior Composition," in *Proc. of the 13th International Symposium on Distributed Autonomous Robotic Systems*, London, 2016, pp. 1–13.

[21] S. Pourmehr, V. M. Monajjemi, R. Vaughan, and G. Mori, "'You two! Take off!': Creating, modifying and commanding groups of robots using face engagement and indirect speech in voice commands," *IEEE International Conference on Intelligent Robots and Systems*, no. i, pp. 137–142, 2013.

[22] D. St-Onge, U. Cote-Allard, K. Glette, B. Gosselin, and G. Beltrame, "Engaging with robotic swarms: Commands from expressive motion," *Transaction on Human-Robot Interactions*, 2019.

[23] F. Fink-Hooijer, "Civil protection and humanitarian aid in the Ebola response: lessons for the humanitarian system from the EU experience," *Humanitarian Exchange*, vol. 64, no. June, pp. 3–5, 2015.

[24] V. S. Varadharajan, B. Adams, and G. Beltrame, "Failure-tolerant connectivity maintenance for robot swarms," *CoRR*, vol. abs/1905.04771, 2019. [Online]. Available: http://arxiv.org/abs/1905.04771

[25] G. Li, D. St-Onge, C. Pinciroli, A. Gasparri, E. Garone, and G. Beltrame, "Decentralized progressive shape formation with robot swarms," *Autonomous Robots*, vol. 43, no. 6, pp. 1505–1521, Aug 2019.

[26] B. Pendleton and M. Goodrich, "Scalable Human Interaction with Robotic Swarms," *AIAA Infotech@Aerospace (I@A) Conference*, pp. 1–13, 2013.

[27] J.-P. Gagné, J. Besser, and U. Lemke, "Behavioral assessment of listening effort using a dual-task paradigm: A review," *Trends in Hearing*, vol. 21, p. 2331216516687287, 2017.

[28] J. Cort, M. Ieee, S. Mart, M. Ieee, T. Karatas, F. Bullo, M. Ieee, J. Cortés, S. Martínez, T. Karataş, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Transactions on Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.

[29] V. Alexandrov, K. Kirik, and A. Kobrin, "Multi-robot Voronoi tessellation based area partitioning algorithm study," *Journal of Behavioral Robotics*, pp. 214–220, 2018.

[30] S. Fortune, "A sweepline algorithm for voronoi diagrams," *Algorithmica*, vol. 2, no. 1, p. 153, Nov 1987.

[31] G. Podevijn, R. O'Grady, N. Mathews, A. Gilles, C. Fantini-Hauwel, and M. Dorigo, "Investigating the effect of increasing robot group sizes on the human psychophysiological state in the context of human-swarm interaction," *Swarm Intelligence*, vol. 10, no. 3, pp. 1–18, 2016.

[32] M. Lewis, "Human Interaction With Multiple Remote Robots," *Reviews of Human Factors and Ergonomics*, vol. 9, no. 1, pp. 131–174, 2013.

[33] D. R. Olsen Jr and S. Bart Wood, "Fan-out : Measuring Human Control of Multiple Robots," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, Vienna, 2004, pp. 231–238.

[34] J. Wang and M. Lewis, "Assessing Coordination Overhead in Control of Robot Teams," *Systems, Man and Cybernetics*, pp. 2645–2649, 2007.

[35] A. Kolling, P. Walker, N. Chakraborty, K. Sycara, and M. Lewis, "Human Interaction with Robot Swarms: A Survey," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 1, pp. 9–26, 2016.

[36] G. Podevijn, R. O'Grady, C. Fantini-Hauwel, and M. Dorigo, "Investigating the effect of the reality gap on the human psychophysiological state in the context of human-swarm interaction," *PeerJ Computer Science*, vol. 2, no. 82, 2016.

[37] J. Heard, C. E. Harriott, and J. A. Adams, "A Survey of Workload Assessment Algorithms," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 434–451, 2018.

[38] R. Buettner, S. Sauer, C. Maier, and A. Eckhardt, "Real-time Prediction of User Performance based on Pupillary Assessment via Eye-Tracking," *AIS Transactions on Human-Computer Interaction*, vol. 10, no. 1, pp. 26–56, 2018.

[39] S. Mathot, "Pupillometry: Psychology, physiology, and function." vol. 1, no. 1, p. 16, 2018.

[40] V. Peysakhovich, F. Dehais, and M. Causse, "Pupil Diameter as a Measure of Cognitive Load during Auditory-visual Interference in a Simple Piloting Task," *Procedia Manufacturing*, vol. 3, pp. 5199–5205, 2015.

[41] J. Schmidtler, K. Bengler, F. Dimeas, and A. Campeau-Lecours, "A Questionnaire for the Evaluation of Physical Assistive Devices (QUEAD)," in *IEEE International Conference on Systems, Man, and Cybernetics*, Banff, 2017.

[42] J. Carifio and R. J. Perla, "Ten Common Misunderstandings, Misconceptions, Persistent Myths and Urban Legends about Likert Scales and Likert Response Formats and their Antidotes," *Journal of Social Sciences*, vol. 3, no. 3, pp. 106–116, 2007.

[43] T. Santini, W. Fuhl, and E. Kasneci, "Calibme: Fast and unsupervised eye tracker calibration for gaze-based pervasive human-computer interaction," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ser. CHI '17. New York, NY, USA: ACM, 2017, pp. 2594–2605.

[44] J. Blamont, "A roadmap to cave dwelling on the moon and mars," *Advances in Space Research*, vol. 54, no. 10, pp. 2140–2149, 2014.