



Control, localization and human interaction with an autonomous lighter-than-air performer

David St-Onge^{a,*}, Pierre-Yves Brèches^b, Inna Sharf^b, Nicolas Reeves^c, Ioannis Rekleitis^e, Patrick Abouzakhm^b, Yogesh Girdhar^d, Adam Harmat^b, Gregory Dudek^d, Philippe Giguère^a

^a Department of Computer Science and Software Engineering, University Laval, Adrien-Pouliot Building, Quebec, Qc, Canada

^b Department of Mechanical Engineering, McGill University, Macdonald Engineering Building, Room 148, Montréal, Qc, Canada

^c School of Design, University of Quebec in Montréal, Montréal, Qc, Canada

^d Centre for Intelligent Machines, McGill University, McConnell Engineering Building, Room 419, Montréal, Qc, Canada

^e Computer Science and Engineering Department, University of South Carolina, Columbia, SC, USA

HIGHLIGHTS

- Presents an art-science-technology project for HRI research in the artistic realm.
- Details the dynamics, with effects of its added mass and vision-based localization.
- Experiments with controlled trajectories of the blimp for evoking emotions.
- Denes a sonar-based sphere of intimacy and a vision-based perplexing features scanner.
- Uses of electromyography sensors to interpret performers moods.

ARTICLE INFO

Article history:

Received 6 July 2015

Received in revised form 7 September 2016

Accepted 17 October 2016

Available online 24 October 2016

Keywords:

Robotic blimp

Robotic art

Human-robot interaction

Dynamic modeling

Mobile robotics

Airship

Theater

ABSTRACT

Due to the recent technological progress, Human-Robot Interaction (HRI) has become a major field of research in both engineering and artistic realms, particularly so in the last decade. The mainstream interests are, however, extremely diverse: challenges are continuously shifting, the evolution of robot skills, as well as the advances in methods for understanding their environment radically impact the design and implementation of research prototypes. When directly deployed in a public installation or artistic performances, robots help foster the next level of understanding in HRI. To this effect, this paper presents a successful interdisciplinary art-science-technology project, the *Aerostabiles*, leading to a new way of conducting HRI research. The project consists of developing a mechatronic, intelligent platform embodied in multiple geometric blimps – cubes – that hover and move in the air. The artistic context of this project required a number of advances in engineering on the aspects of localization and control systems, flight dynamics, as well as interaction strategies, and their evolution through periods of collective activities called “research-creation residencies”. These events involve artists, engineers, and performers working in close collaboration, sometimes, over several weeks at a time. They generate fruitful exchanges between all researchers, but most of all, they present a unique and creative way to direct and focus the robotics development. This paper represents an overview of the technical contributions from a range of expertise through the artistic drive of the *Aerostabiles* project.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

The realm of robotics has always required a wide variety of expertise to ensure that the results fully correspond to the objectives and intentions of the designers. Unfortunately, the robotic

development most frequently occurs in specialized robotics research facilities and involves homogeneous groups of highly skilled technical experts and engineers. Although multidisciplinary teams slowly started to appear in universities in the last decade, involving engineers from different disciplines as well as psychologists, such groups are often missing the people that are seen as the experts in the fields of creativity, perception, and sense-crafting—the ubiquitous artist.

* Corresponding author.

E-mail address: david.st-onge.2@ulaval.ca (D. St-Onge).



Fig. 1. Aychel Szot, performer, with two Tryphons, cubic autonomous blimps.

Simultaneously, several major robotic artworks were created by visionary artists such as White, Roussi, Shanon, Vorn, and Moura [1]. Norman White's "Facing Out Laying Low"¹ was among the first interactive microcontroller based artistic installations embedded with a basic "intelligence". It reacted to the surrounding light and sounds and expressed its enthusiasm by whistling. This art installation came a decade after Shanon's reactive "Squat robot",² which was triggered by visitors when touching a plant [2], and a decade before Gilles Roussi's "Bons robots",³ a seven-meter automaton meant to inspire fear in the audience. Opposed to Roussi's creation, White's later work "Helpless Robot"⁴ was deliberately non-anthropomorphic, but managed to inspire empathy in the visitors. Seeking a similar impact, in the mid-90's Bill Vorn's "Hysterical Machines"⁵ were insect-like robots that reacted violently to the presence of visitors, but their reactions were as intense as useless, which again inspired feelings of compassion and sorrow. Most recently, Leonel Moura's "RAP"⁶ is comprised of entirely autonomous painting robots creating artworks from basic knowledge of symbols, text and aesthetic principles programmed by the artist. Instead of task-focused robots, these machines question the very notion of the meaning and significance of robot behaviors. The visitors' interpretations of these behaviors is a powerful interaction vector, but its comprehension is still mostly orthogonal to roboticians.

Central among this prestigious group of key robotic artworks stands the *Aerostabiles*⁷ project, originally conceived by artist, architect and professor Nicolas Reeves in 2002 [3]. Borrowing from an architectural background, the concept is to translate the age-old myth of a heavy mass freed from the laws of gravity, a myth

that many predecessors tried to materialize through various strategies [4]. The basic building block, a cubic-shaped brick embodied in a cubic helium blimp, was selected and imagined to levitate in such a way that all positioning, flying and controlling processes would disappear only to leave a meditative and contemplative visual artwork as seen in Fig. 1. A first basic motorized version, called *Mascarillons*, based on a 1.80-meter cubic basswood structure, flew during two weeks for the first, fully public installation at the Quebec Museum of Civilization in Quebec City, Canada [5]. A first narrative performance subsequently occurred at the Montréal Centre for Science, telling the story of a 1.60-meter cube (Nestor) behaving like a large animal being tamed by an expert circus tamer. These early prototypes were not reliable and had limited capabilities, but the interest for the project was progressively growing. A larger and more robust version, the Tryphon, was created, made of a 2.25-meter carbon fiber structure with a helium filled bladder, with an increased autonomy so as to allow for long-term installations. The Tryphons were shown in a number of international events around the world [6,5], namely in the *Summer of Dance* in Paris, hovering over a dancing crowd with video projections on their square faces and in the *Science as Suspense* event in Moscow. In each of their appearances, the behavior and environment of the Tryphons were different and required specific technical developments. In order to enhance the performance possibilities of the floating vehicles, to increase their autonomy and to limit the adaptation and calibration time required for each new environment, an interdisciplinary team project was needed. In 2012, three robotic engineering labs (two from McGill and one from Laval University) started working together with the artists on the specific challenges of localization, precise control and interactivity of the Tryphons. This paper presents the results of that collaboration. It is structured in such a way that each technical contribution is presented in relation with the requirements and questions that arose from the artistic application of this unique project. Building on the control and localization solutions developed, Section 6 ends the paper with the results of new human–robot interaction experiments and exploration of novel devices in the context of the *Aerostabiles* project.

2. Infrastructure of the Aerostabiles

As the most recent generation of the *Aerostabiles* prototypes, the Tryphons incorporate the cumulative improvements developed through their evolution during eight years of artistic performances.

¹ White, N. *Facing Out Laying Low*. 1977. Plexiglass, motors, electronics. Collection Norman White. Ottawa.

² Shanon, T. *Squat robot*. 1966. Live plant, electronic sensors, painted metal, motor, rollers. John Kingsley Shanon Collection. Caledonia.

³ Roussi G. *Bon robots*. 1984. steel, glass, electronics. Cité des sciences et de l'industrie. Paris.

⁴ White N. *Helpless Robot*. 1987. Plywood, steel, proximity sensors modified 80386 computer. Agnes Etherington Art Centre. Ottawa.

⁵ Vorn B. *Hysterical Machines*. 2006. Aluminium, proximity sensors, motors. Bill Vorn Collection. Montréal.

⁶ Moura, L. *RAP*. 2014. Aluminium, wheels, painting pen, electronics. Leonel Moura Collection. Lisbon.

⁷ The project name is an original idea of Nicolas Reeves, inspired by Alexander Calder's *Aerostabiles*.

Table 1

Weight of different elements for basic configuration of the Tryphon as in Fig. 2.

	Weight (g)	Remaining payload (g)
Helium buoyancy	N/A	9600
Bladder	3000	6600
Structure & ducts	4000	2600
Cables & hubs	682	1918
Motors	800	1118
ESC	256	862
Batteries	780	82
Gumstix boards	66	16

In their current state, they serve as a research and exploration platform that allows multidisciplinary research–creation residencies (in situ exploration and development) involving more than 25 people, during which teams from diverse disciplines meet and exchange ideas to push the boundary of the flying cubes' potential for art, science and engineering.

2.1. Tryphons hardware

In order to be usable for long-lasting performances and installations, the Tryphons need to carry as many on-board batteries as possible and to be equipped with efficient actuators. Since the *Aerostabiles* are lighter than air, their maximum payload depends directly on the volume of helium they contain. Each component was carefully selected and designed to optimize its weight to functionality ratio. The weight breakdown of the available payload is summarized in Table 1 for the configuration of Fig. 2. In order to give a global overview of the project, the key elements of the design are recalled here from [3].

2.1.1. Structure

The size of Tryphon blimps, although preferred to be larger, was limited by the fact that Tryphons are frequently shipped by air cargo around the world, constraining the size of their trusses to 2.25 m so as not to exceed the threshold for over-sized cargo transport. The structure fabrication is simplified by using carbon-fiber rods, tubes and strips glued together in the right configuration with 3D printed joints. To optimize and accelerate the assembly on each deployment, no mechanical attachments are required for the twelve trusses to form the cube: the trusses slide into each other and are attached to the cubic helium bladder, whose tension when inflated stabilizes and secures the whole structure together. The structure also supports all the mechatronics components described in the following sections.

2.1.2. Actuators

The Tryphon may be operated for days in a performance setup and thus, to actuate this large vehicle whose flat faces act like sails, the design requires high-efficiency reliable propellers. The Tryphons are currently equipped with motorized carbon fiber ducted fans actuated by brushless DC motors. The robustness of the motors and the carbon-made ducts are well suited for extensive use over the years, and are reliable enough for presentation in front of large audiences. Like all devices on the Tryphons, the fans are attached to the structure with custom-made resin clips: they can be moved or replaced quickly (<60s). For safety and reliability, redundancy is an important feature of the design. Even if six motors are in theory sufficient to control the six degrees of freedom of the blimp, more are usually installed. This also substantially increases the total available thrust, and helps to ensure symmetrical application of the thrust to the Tryphon. Two different thruster configurations have been experimented with: one with 12 thrusters and one with 8, the latter schematically illustrated in Fig. 2.




Fig. 2. Embedded hardware –basic configuration. All components are mounted on the trusses (edges).

Each thruster motor can draw up to 20 A of electrical current from the onboard 2-cell LiPo battery packs (7.4 V). The fans are commanded in percentage of the maximum velocity with an 8-bit resolution scale (0–255) and a second argument for the direction of rotation sent through the I²C bus to their Electronic Speed Controllers (ESCs). Bench tests have been conducted with the thruster assembly, as illustrated in Fig. 3, to produce thrust vs. input velocity command calibration. In these tests, thrust values were measured with a force–torque sensor mounted at the bottom of the vertical mast supporting the whole truss assembly. The latter includes the propeller, its ESC and the polycarbonate ducts extending from the fans on each side to the corners of the cube; these ducts are intended to keep the airflow as laminar as possible (see Fig. 4). For safe usage over the long term, the maximum velocity command is limited to 75%: over this limited range, experiments show a near linear relation between the resulting thrust and the velocity command (see Fig. 5). The thruster performance in such a setup, which closely mimics the actual thruster arrangement on Tryphon, can hardly be compared to the nominal performance of the propellers as per manufacturer specifications, because of the lower operating voltage, the presence of the truss and tubes, and the reduced maximum propeller velocity. Note that the relationship shown in Fig. 5 is slightly asymmetric due to the thruster mechanical construction.

A key component when dealing with brushless motors is the Electronic Speed Controller (ESC). Selection of this device,⁸ which determines the performance and reliability of the actuator, was based on the accessibility of software sources and hardware schematics, in order to fit the constantly evolving technical needs of the project. Functionality was added to allow forward and backward rotation selection (not needed for most flight ESC), and to send battery monitoring information. Robustness to battery voltage fluctuations at high-load was improved by separating the ESC micro-controller power from the motor battery. To get a symmetric distribution of the weight, the batteries are located in each corner of the cubic structure. At a full twelve fans configuration, 6 batteries power the actuators while additional two serve to power the sensors and on-board computer. The eight fans configuration

⁸ BL-Ctrlv1.2 from <http://www.mikrokopter.de>.




Fig. 3. View of the bench setup to test the ducted fans performances.




Fig. 4. View of a carbon-fiber fan from output of plastic duct at the corner of the cubic structure.




Fig. 5. Relation between output thrust and input velocity command for the Tryphon's propellers.

uses four batteries for the actuators and only one for its logic component, located on the bottom layer of the structure, symmetrical to the Gumstix board. In all configurations, the thrusters share batteries, in order to limit the use of the available payload.

2.1.3. On-board computer

For internal communication, the *Aerostabiles* prototypes rely mostly on the I²C bus. This protocol was selected for its simplicity of implementation and the wide range of sensors and actuators

already compliant with it. Most microprocessors have I²C protocols implemented, but the Tryphons also require a wireless interface, USB ports and sufficient processing power for both control algorithms and video processing, all while complying with the payload constraints. Amongst the devices on the market meeting these requirements and that are manufactured by a stable, well-established company, the GumstickTM computer-on-module product line was selected. Their Overo Firestorm⁹ built with an ARM Cortex-A8 microprocessor with video accelerator fully supports Linaro.¹⁰ It is complemented with the Robovero¹¹ expansion board that includes an Inertial Measurement Unit (IMU) and multiple I/O pins.

2.1.4. Sensors, bus and hubs

Following the aforementioned need to adapt the sensor configuration to different performance spaces, the electronics on the Tryphons has to be modular. Custom hubs allow for easy connection and reconfiguration of the sensors plugged in the robots with robust RJ9 connectors.

Over the years, many different sensor sets have been experimented with. Multiple sonars (ultrasound sensors) were used in the initial deployments of the aerobots as the only means of localization. This sensor modality proved to be effective on other robotic platforms [7], even for indoor blimps [8]. Their use on Tryphons, however, was only reliable in spaces where reference objects are at less than six meters from the aerobot in at least three directions. This restricted the movements of the Tryphon to a room corner. The orientation (heading) was then either given by computing the angle from the distance values of two separated sonars on the same face or by a gravity compensated compass. To allow for more flexible and robust localization, a recent upgrade involved the integration of a full 9 degrees-of-freedom IMU from the Robovero expansion board and a number of USB cameras. For ground truth and to help with some aspects of development, an external motion capture system¹² and an external 2D laser scanner (SICK)¹³ were also sometimes used. The laser has a range of up to 50 m but its field of view is only 120°, thus limiting its tracking capabilities. The vision-based localization strategy currently implemented on Tryphons is described in Section 4.

Sonars have also been used as interaction devices with the audience and performers, as for example, in experiments where artists stimulate the sonars while the aerobot hovers in a corner, thus provoking perturbations to which the Tryphon reacts. Other sensors have been experimented with in the more recent residencies: light sensors and cameras on-board, microphones and electromyogram on the performers. These interactive systems will be discussed in Section 6.

2.2. The need for residencies work—merging of academic cultures

Amongst the tools and strategies developed by researchers and artists participating in the Tryphons project, the most unique is the organization of regular collaborative activities based on the model of art residencies. These events, scheduled twice a year, are structured around the framework of engineering software and hardware integration workshops, which occur conjointly with artistic practices and creations. During these research–creation residencies, engineers and artists from different disciplines work together over intensely collaborative periods lasting from one to

⁹ <https://store.gumstix.com/index.php/products/622/>.

¹⁰ A Debian-based linux distribution project for ARM.

¹¹ <http://robovero.org/>.

¹² Vicon from <http://www.vicon.com/>.

¹³ LMS200 from <https://www.sick.com/>.




Fig. 6. Tryphon body-fixed frame (CM-xyz), propeller nominal thrust directions F_i , Euler rotation axis (ψ, θ, ϕ) and inertial reference frame (XYZ).

three weeks, in spaces large enough to allow for the flight of at least two Tryphons.

Abundant questions and discussions emerge from these residencies, leading to new investigations and development foci both in art and engineering. Whenever possible, residencies take place in public spaces, allowing direct contact with diverse audiences, and with inspiring media artists and creators—researchers. The specific constraints of out-of-the-lab environments raise new challenges for engineers, while the encounters between different academic cultures influence the development priorities of the *Aerostabiles* project. For example, the modularity of the mechatronic design and the software architecture of the Tryphons (to be described in Section 5.1), which are required to enable quick behavior and interaction design for the artists, both of these needs emerged over the course of residencies. Another example of the need identified through the residencies is for localization backup solutions for when the artists wish to change the lightning or physical configuration of the environment. Finally, system robustness is always challenged in the harsh contexts of residencies, from that of the truss structure to the reliability of the wireless network communications.

The main common point between artists and scientists is the fact that both are trying, each with their own means, to approach the frontier of the unknown: to reach a point beyond which nothing can yet be said, and from which territories that nobody has ever explored can be observed and described, inventing a new model of artistic expression or a new scientific model if needed. It is the *mode* of approach that distinguishes science from arts. Before stating anything new, a scientist must undertake an exhaustive recapitulation of everything that has been said in the specific research field. This long and often tedious process is the *sine qua non* condition for any worthwhile discovery; and the models that the scientist uses to do so are often very remote from the reality of everyday life. By contrast, the artist reaches the unknown by observing first that the deepest mysteries can be hidden very close to us, and even inside the most ordinary things.

It is through these considerations that the project team manages to find elements of a common language which allows to share methodologies and results in a way that is fruitful for everyone. A vibrant instance of this duality between the approaches and aims of art and science is the understanding of the special dynamics of the Tryphons. Indeed, the engineers apply Newtonian laws to develop a mathematical model and then adapt the model parameters to the reality of experimental data, as will be shown in the next section, while the performers dance and play for several days with a Tryphon whose devices are turned off. Through this extensive experimentation, they manage to find an intuitive way to understand the specific rhythm and pace, i.e., the dynamics of the Tryphon movements. This allows them to grasp the complexity

of the processes that are necessary to precisely control the automaton, and the level of energy required even for seemingly simple movements, such as a brisk turn or a deceleration.

3. Tryphon dynamics

Over the course of Tryphon development as part of the *Aerostabiles* project, it became evident that to develop effective controllers for the blimp in order to enable aesthetic motion in large spaces, as well as interaction with performers and audience during artistic performances, a high-fidelity dynamics simulator of the blimps was required. In the broad category of aerial robotic vehicles or unmanned aerial vehicles, Tryphon, as a helium-inflated blimp, falls into a class of lighter-than-air vehicles. As such, its dynamics is governed by standard rigid-body equations of motion, but with notable additions: the force of buoyancy and the added-mass effect. The dynamics model formulated for Tryphon is summarized here.

3.1. Equations of motion

3.1.1. Reference frames and kinematics

The six generalized coordinates $x, y, z, \phi, \theta, \psi$ are used to represent the Tryphon pose in the inertial frame (XYZ, see Fig. 6). The former three coordinates represent components of the position of the center of mass (CM) of Tryphon in the inertial frame and the latter three are Euler angles describing the orientation of the Tryphon body-fixed frame (CM-xyz), following the flight dynamics convention ZYX as in the work of Bishop [9]. Thus, the transformation between the angular velocity components and Euler rates is given by:

$$\omega = S\dot{\phi} \text{ where } S = \begin{bmatrix} 1 & 0 & -\sin \theta \\ 0 & \cos \phi & \cos \theta \sin \phi \\ 0 & -\sin \phi & \cos \theta \cos \phi \end{bmatrix}, \quad (1)$$

with $\omega = [p, q, r]^T$ are the body frame angular velocities and the Euler rates are $\dot{\phi} = [\dot{\phi}, \dot{\theta}, \dot{\psi}]^T$. The rotation matrix from the inertial frame into the Tryphon (body) frame will be denoted with R .

3.1.2. Dynamics model

Having defined the necessary reference frames, the equations of motion are derived by employing the Newton–Euler formulation, with the translational equations expressed in the inertial reference frame, while the rotational equations of motion, taken about the center of mass of the Tryphon, are formulated in Tryphon body-fixed frame. In combined matrix form and following the notation in [10], these equations are written as:

$$\dot{M}\mathbf{V} = \tau_I + \tau_V + \tau_G + \tau_C \quad (2)$$

where $\mathbf{V} = [\mathbf{v}^T, \boldsymbol{\omega}^T]^T$ is the vector of generalized velocities, $\mathbf{v} = [\dot{x}, \dot{y}, \dot{z}]^T$ is the velocity of the center of mass of the Tryphon blimp expressed in the inertial frame and M is the generalized mass matrix and is defined as:

$$M = \begin{bmatrix} m\mathbf{I}_{3 \times 3} + \mathbf{A}_m & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{J} + \mathbf{A}_j \end{bmatrix} \quad (3)$$

with m and $\mathbf{J}$ being, respectively, the mass and the matrix of inertia of Tryphon; accurate estimates of these inertial parameters are available from a detailed CAD model and the weighting of each components. The matrices $\mathbf{A}_m$ and $\mathbf{A}_j$ are, respectively, the added-mass matrix and the added-inertia matrix and are defined as in the work of Korotkin [11]:

$$\mathbf{A}_m = A_m \mathbf{I}_3 = (k\rho_{air}V) \mathbf{I}_3 \quad (4)$$

$$\begin{aligned} \mathbf{A}_j &= A_j \mathbf{I}_3 - A_m \mathbf{r}_B^\times \mathbf{r}_B^\times \\ &= \left(k' \rho_{air} \frac{Vl^2}{6} \right) \mathbf{I}_3 - (k\rho_{air}V) \mathbf{r}_B^\times \mathbf{r}_B^\times \end{aligned} \quad (5)$$

where V is the volume of the Tryphon, l is the side length of the Tryphon cube, ρ_{air} is the density of air, k is the added-mass coefficient, k' is the added-inertia coefficient, $\mathbf{r}_B^\times$ is the cross product matrix associated with the position vector of the geometric centroid of the Tryphon blimp (center of buoyancy) with respect to the center of mass, expressed in the body-fixed frame as $\mathbf{r}_B = [0, 0, l_b]^T$. Note that the added mass is represented by a scalar matrix due to symmetry of the cube. Also noted is that Eq. (5) embodies the application of parallel-axis theorem to transform the added inertia from the center of buoyancy to the center of mass for use in dynamics equation (2). The computation of the added-mass and added-inertia coefficients will be discussed in more detail in Section 3.1.3.

The right-hand side of equations of motion is comprised of four generalized force terms, as follows. The vector $\boldsymbol{\tau}_I$ contains the Centrifugal moments due to the rotating frame used to formulate the rotational equation of motion and is defined as:

$$\boldsymbol{\tau}_I = \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ -\boldsymbol{\omega}^\times \mathbf{J} \boldsymbol{\omega} - \boldsymbol{\omega}^\times \mathbf{A}_j \boldsymbol{\omega} \end{bmatrix}. \quad (6)$$

Due to its non-aerodynamic shape, the viscous effects of air on the Tryphon blimp are not negligible and are described by the vector $\boldsymbol{\tau}_V$ defined as:

$$\boldsymbol{\tau}_V = \begin{bmatrix} \mathbf{F}_D \\ \mathbf{r}_B^\times \mathbf{R} \mathbf{F}_D + \mathbf{M}_D \end{bmatrix} \quad (7)$$

where $\mathbf{F}_D = -\frac{1}{2} \rho_{air} C_d A |\mathbf{v}| \mathbf{v}$ is the aerodynamic drag force, applied at the geometric center of the cube. The aerodynamic drag moment $\mathbf{M}_D$ is obtained based on the derivation in [12], simplified for a cubic body as:

$$\mathbf{M}_D = -\frac{1}{32} \rho_{air} C_d l^5 \begin{bmatrix} |p|p \\ |q|q \\ |r|r \end{bmatrix}. \quad (8)$$

Additional parameters employed in the above equations are the projected area A of the Tryphon blimp normal to $\mathbf{v}$ and the drag coefficient C_d of the Tryphon blimp [13].

The forces and moments created by gravity and buoyancy are combined in the vector $\boldsymbol{\tau}_G$ defined as:

$$\boldsymbol{\tau}_G = \begin{bmatrix} \alpha \mathbf{F}_G \\ \mathbf{r}_B^\times \mathbf{R} \mathbf{F}_B \end{bmatrix} \quad (9)$$

where $\mathbf{F}_G$ and $\mathbf{F}_B$ are, respectively, the gravity and buoyancy forces and α is the percentage difference between the buoyancy and the gravity force. In the ideal case of $\alpha = 0$, Tryphon is in the state of neutral buoyancy and remains stationary in mid-air without application of any thrust.

Finally, the resultant of the forces and moments generated by the eight propellers actuating the blimp is represented by the vector $\boldsymbol{\tau}_C = [\mathbf{F}_C^T, \mathbf{M}_C^T]^T$ where

$$\mathbf{F}_C = \mathbf{R}^T \mathbf{L}_1 \mathbf{F}_p \quad (10)$$

$$\mathbf{M}_C = \mathbf{L}_2 \mathbf{F}_p \quad (11)$$

and the transformation matrices required are:

$$\mathbf{L}_1 = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix} \quad (12)$$

$$\mathbf{L}_2 = \begin{bmatrix} 0 & 0 & l_z & l_z & -l_y & -l_y & l_y & l_y \\ l_z & l_z & 0 & 0 & -l_x & l_x & l_x & -l_x \\ l_y & -l_y & l_x & -l_x & 0 & 0 & 0 & 0 \end{bmatrix} \quad (13)$$

with $l_z = \frac{l}{2} - l_b$ and $l_x = l_y = \frac{l}{2}$. The column vector $\mathbf{F}_p$ contains the thrusts generated by the eight propellers mounted on the Tryphon blimp (see Fig. 6). The estimation of the added-mass and added-inertia coefficients is non-trivial and is specifically addressed in the following section.

Table 2

Added-mass coefficients for Tryphon as function of n_p .

n_p	600	2400	3750	6144	9600
k	0.6588	0.6471	0.6450	0.6433	0.6420
k'	0.2618	0.2480	0.2452	0.2428	0.2411




Fig. 7. Added-mass coefficients estimated for different thrust levels and position measurement sensors in linear motion experiments; $k = 1.04$, $\sigma = 0.21$.

3.1.3. Estimation of added mass coefficients

The added mass effect is an important aspect of modeling a lighter-than-air flying object and is particularly critical for the dynamics of Tryphon because of its non-aerodynamic cubic shape. This effect originates from the mass of fluid moved by the body under consideration while accelerating. Since its discovery, various methods have been developed to extract the added-mass coefficients [14–16].

The boundary element method (BEM) algorithm based on the work of Ghassemi and Yari [14] was used on a Tryphon cube meshed with a progressively increasing number of patches n_p and the results for the added-mass and added-inertia coefficients are summarized in Table 2. As can be seen, the values of k (added-mass coefficient) and k' (added-inertia coefficient) are converging and are coherent with the lower bounds given by the theoretical coefficients of a sphere, which are 0.5 and zero, respectively. The value of k computed here is very close to the added-mass coefficient of 0.638 estimated for a cube in [17] with the variant of the BEM.

In addition, an experimental investigation of the added mass of Tryphon was undertaken to corroborate the numerical results obtained with the BEM. The added-mass coefficient was estimated from a series of constant-thrust rectilinear translation tests, which were realized by thrusting the blimp in one direction using different percentages of the maximal thrust ([0.35, 0.50, 0.60, 0.85]). Two types of sensors were used to measure the distance during these translational tests: sonars on the Tryphon and the ground-fixed 2D laser scanner, for its larger measurement range. The measurements of thrust and distance traveled were fitted with the analytical solution of the rectilinear equation of motion:

$$\dot{V}_x = -\frac{1}{2} \rho_{air} C_d A v_x^2 + F_x, \quad (14)$$

where F_x is the thrust computed from the thruster input-output map (Fig. 5) and $\dot{V}_x$, v_x are the acceleration and velocity of the blimp in the direction of translation, estimated from the distance measurements.

The experimental results at different percentages of maximum thrust and for the two measuring modalities are displayed in Fig. 7. The average estimate for the added-mass coefficient k is 1.04 which is 63% higher than the value computed with the BEM. This can be partly attributed to the imperfect shape of the Tryphon blimp:

the external structure and the bulging of the flexible faces are not taken into consideration in the perfect cubic shape modeled by the BEM. Another contributing factor is likely to be due to the floor effect: in the translational experiments, the Tryphon was flown approximately 1.5 m above the floor. In a previous investigation conducted by Vorobjov [15] with a parallelepiped submerged in water between two parallel walls at a distance of 1.5 m from each, the authors observed an increase of approximately 80% in the added mass, compared to the result in an infinite fluid.

The added-inertia coefficient was estimated by carrying out free-oscillation tests and using a basic relationship between the measured period of oscillation and the blimp inertia. These tests produced the average added-inertia coefficient estimate of $k' = 0.25$, with standard deviation of $\sigma = 0.09$, which is in excellent agreement with the value of 0.24 computed with the BEM.

3.2. Simulation environment

Since experiments with the Tryphons are limited to the residency periods, an accurate simulation environment is necessary in order to develop and evaluate the controllers and to test different interaction behaviors prior to residency trials. Two simulation models were developed: a Simulink model, chosen for the simplicity of implementation it entails, and a Gazebo model, for its compatibility with the Robotic Operating System (ROS [18]), the software platform employed on the Tryphons (see Section 5.1).

The Simulink modeling environment was first investigated for the dynamics model, control and force distribution component blocks to represent “ideal” conditions. Later on, feedback noise and thruster saturation blocks were added to achieve a more realistic simulation. The results can be either directly plotted in Simulink or sent to a MATLAB workspace, providing fast turnaround for tuning controller gains, adjusting noise and saturation levels and evaluating different control law designs.

The Gazebo simulation software uses OGRE engine to render a 3D virtual world in which the Tryphon can be emulated, providing more powerful visualization tools. This software is compatible with ROS and allows, with minimal modifications, the same ROS nodes to be used on both the emulated and real Tryphon. Furthermore, sensors can be added in Gazebo with an included noise model; this has proven particularly useful for incorporating feedback from our vision based localization algorithm (MCPTAM) by streaming images from virtual cameras. Thus, the Gazebo environment allows seamless modeling and integration of sensors, thereby improving the fidelity of the simulator. In addition, not only can Gazebo be used for controller development and system performance analysis, it also allows testing and debugging of all custom software before its deployment on the real system. Fig. 8 illustrates the Gazebo simulation environment, modeled after Mount Royal Chalet in Montreal where a recent residency took place; the virtual cameras take images of the virtual walls in order to localize the aerobot using MCPTAM. Recently, the Gazebo simulator was expanded by integrating a second Tryphon to develop controllers for docking and collaborative behaviors.

4. Localization

To create scripted or autonomous trajectories, the aerobots need a reliable localization system. Such a system must be quickly and easily set up and calibrated and, function in a variety of spaces where the shows take place, for any potential scenario or narrative in which the Tryphon can be involved. To ensure success of an artistic event, the system needs to be reliable and robust, which can be achieved by merging several localization strategies. Current research on unmanned aerial vehicles control in GPS denied environments relies heavily on the use of motion capture equipment,

such as Vicon, to track the pose of the robots [19,20]. The cost and time of deploying such a device are not compatible with the constraints of art events. Nevertheless, such efforts have been made previously for the Tryphon project [21] and the 2D laser scanner noted earlier has been integrated in some of the recent residencies. Indeed, this scanner can be set up in less than an hour and does not require extensive calibration since the custom tracking software only needs to follow a line in a cross-section.

From an artistic perspective, the robots need to have some degree of autonomy so as not to rely on prior room preparation or external sensing devices. Therefore, strong emphasis has been placed in the *Aerostabiles* project on on-board localization methods. The IMU on-board the Tryphon gives accurate measurements of the aerobot roll and pitch angles with its accelerometer and, the magnetometer, often paired with a compass, measure the yaw of the blimp, but the absolute positions x , y and z are more difficult to obtain. As previously stated, the first set of sensors to be used on-board for this purpose were sonars with a range up to 6 m. These light-weight and inexpensive devices, however, have limited capabilities for navigation and control. While some implementation of probabilistic localization using sonars have been successful [8], these sensors have mainly been used for obstacle detection and stabilization in the *Aerostabiles* project. As sonars intrinsically produce relative localization specific to objects in the environment (e.g., walls, ceilings) they were found more relevant for use as an analog interaction interface, as will be discussed in Section 6. The novel vision-based localization techniques developed for this project are the relative pose determination from a mutually-observing pair of cameras, *You see Me, I see You* and the global localization and mapping solution using multiple cameras, *MCPTAM*.

4.1. You see Me, I see You

As the Tryphons are meant to operate in small groups, collaborative-type localization approaches can be leveraged. These approaches necessitate the capability to perform relative localization between vehicles, to be able to fuse individual pose estimates together in order to improve overall precision. Within the scope of this project, we developed such a relative localization approach in 6DoF, proposed by Dugas et al. [22]. It is itself an extension of a 3DoF algorithm proposed earlier by Giguere et al. [23]. They are only based on angular measurements between landmarks on the vehicles themselves, estimated from images taken from onboard cameras. As such, it respects the important criteria of onboard localization.

As these approaches do not use any range (distance) measurements, they are often referred in the literature as *bearing-only* methods. Our solution is inexpensive, accurate at short range, low-weight, and does not require external hardware. Moreover, it scales well to a large number of robots. Consequently, it is uniquely suited to this project.

The problem of relative pose estimation can be formalized as follow. We have two vehicles, *Tryphon_A* and *Tryphon_B*, for which we seek to estimate the relative distance and orientation; see Fig. 9 for the relationship between the coordinate systems of the two robots. Two images, I_A taken by *Tryphon_A* and I_B taken by *Tryphon_B*, are recorded at the same time. From these two images, a number of angles are estimated. Firstly, we estimate the two angles α and β :

- from image I_B : $\alpha = \widehat{L_A C_B R_A}$, which is the angle formed between the markers of *Tryphon_A* and camera C_B ;
- from image I_A : β , the angle between the optical axis of C_A and the ray passing through the origins of C_A and C_B , where the position of C_B is approximated from the location of L_B and R_B .




Fig. 8. Virtual environment and Tryphon in Gazebo. This render is similar to the environment of the Chalet Mont-Royal, on the top of Mont-Royal in Montréal city, Canada.




Fig. 9. The relative localization problem in 3-D, for the two Tryphons A and B operating in 6 DoF. The red and green dashed lines represent the ray between the markers on Tryphon A (L_A and R_A) and the center of projection C_B . The α angle is computed between these rays. The blue dashed line represents the estimated ray between C_B and the center of projection C_A . The angle β is computed between the optical axis z_A of camera A and this ray. On the Tryphons, the markers L_i and R_i are located near the corners of the cubic structure and the cameras C_i in the middle of a strut. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

With these two measured angles α and β and the a-priori known span distance d between the markers on a vehicle, a closed-form approximate solution yields the distance $l = |C_A C_B|$ between the cameras, and hence the vehicles [23]:

$$l = \frac{d}{2 \sin \alpha} \left(\cos \alpha \cos \beta + \sqrt{1 - \cos^2 \alpha \sin^2 \beta} \right). \quad (15)$$

As shown in Fig. 10, tests performed in 2D on an initial setup, with a distance between the landmarks of $d = 0.759$ m, gave an uncertainty on l just above 0.04 m. These results were for ranges of 7–14 m, which would be at the higher end for a typical artistic performance. Importantly, this represents an upper bound on the distance l error. It has been demonstrated in the work of Giguère et al. [23] that it is proportional to the landmark span d . The precision for a Tryphon system would be approximately 3 times better, since $d \approx 2.25$ m in this case.

Note also that this noisy distance estimate l , which is the most important source of uncertainty in our approach, can be improved by performing the same computation described in Eq. (15) a second

time, but by extracting α from I_A and β from I_B , and averaging the computed l s.

The relative $[x, y, z]$ position between the cameras is then found by extending the vector going from C_A to the location of *Tryphon*_B in the image frame to a length of exactly l . Sufficient information is contained in the two images I_A and I_B to recover uniquely both this direction vector and the relative orientation between the two vehicles. The latter corresponds to a rotation matrix that:

- aligns the perceived plane containing C_B , L_A and R_A with the perceived plane containing C_A , its right marker R_A and the other camera C_B ; and
- aligns the perceived vectors $\overrightarrow{C_A C_B}$ in I_A and $\overrightarrow{C_B C_A}$ in I_B in opposite directions.

Since cameras can be considered as excellent protractors, the angular error of this rotation matrix will be comparable to the angular error of the camera, which is limited by the pixel density and the landmark location estimation in the image. This is in contrast to the work of Faessler et al. [24], where the orientation error increased as a function of the distance, as a result of computing relative positions and angles from a single image. Indeed, they reach an average of 3° at $l = 5$ m, compared to an average error of 1.36° at $l = 12$ to 15 m with our approach [25], independent of the distance l .

When dealing with multiple cubic aerobots, the proposed relative localization system will be essential for docking them together, and for reaching a charging station autonomously. Four Tryphons able to assemble in a large space following the directives of artists and architects become a low resolution 3-dimensional floating structure printer: a new way of exploring simple shape design. For global localization, the proposed relative localization scheme can be fused with another method that relies on exteroceptive measurements, such as the MCPTAM described in the following section.

4.2. Multi-camera localization

Localization of Tryphon in its environment is vital if the blimp is ultimately to have the capability for autonomous movement and trajectory following in response to or, in cooperation with a




Fig. 10. Comparison between the estimated relative pose (simplified in 2D) and ground truth established via a laser scanner system, establishing the precision of the distance l for 2 simple trajectories. Experiments were conducted on an early prototype, with a landmark span distance of $d = 0.759$ m, significantly smaller than for a Tryphon system.




Fig. 11. SICK vs. MCPTAM position of Tryphon and distance error between two positions.

human artist. Recently, Harmat et al. [26,27] proposed a multi-camera parallel tracking and mapping (MCPTAM) solution to resolve the localization problem for small unmanned aerial vehicles (UAVs), such as quadrotors. MCPTAM, amongst its advantages, provides much flexibility in terms of positioning, number, type and triggering of the on-board cameras, which make it suitable for implementation on the Tryphon.

MCPTAM is built on the successful monocular parallel tracking and mapping algorithm, combined with the results achieved in the field of multi-camera egomotion estimation. In its full implementation, MCPTAM is a vision-based SLAM solution which provides the global pose of the vehicle by tracking the cluster of rigidly-mounted cameras on board the vehicle, while simultaneously building up the map of the vehicle's environment. As reported in the work of Harmat et al. [27] and of Tribou et al. [28], MCPTAM has been tested in a wide range of indoor and outdoor scenarios and has proven to perform particularly well in urban/man-made environments where planarity of features can be relied on. Furthermore, MCPTAM has been demonstrated to provide very accurate pose estimation of a quadrotor vehicle in flight. This solution gives

the absolute position and orientation of the Tryphon relative to the vision world defined by MCPTAM at initialization.

4.3. MCPTAM implementation on Tryphon

The full implementation of MCPTAM on Tryphon, where the images from several cameras are streamed to the ground station for execution of tracking and mapping threads in parallel has proven to be impractical. The Tryphons are meant to be used as entirely autonomous creatures, and therefore, need to have all the systems and control software required for localization and human interaction on-board the vehicle. However in their present configuration, the main on-board computer is a 1 GHz ARM CPU with 512Mb of flash memory while the minimum requirements are 1.6 GHz and 1G of RAM, and thus attempting to run two synchronized cameras leads to a frame rate of 2 Hz. This is insufficient for closed-loop control of the blimp.

However, since the Tryphon is deployed repeatedly in the same indoor environment for the duration of an artistic event, a strong case can be made for using a pre-built map instead of requiring the system to build a map in parallel with its localization. For



Fig. 12. Views of the Tryphon test environment.

these reasons, MCPTAM is currently employed on the Tryphon in its “partially-autonomous” mode of operation as follows. The pre-built map is constructed under human supervision by using a hand-held multi-camera rig connected directly to a ground station; the map can be saved, re-loaded and post-processed with the help of the custom-designed MAPEDITOR. Once the map is judged to be correct by the operator, the Tryphon blimp can be tracked relative to the mapped space by executing the tracking part of MCPTAM only, thereby alleviating both the computational burden and the communication demands on the on-board hardware. Indeed, as was done for the closed-loop control experiments reported in Section 5.2, a single camera can be used for tracking the blimp relative to a pre-built map.

4.4. MCPTAM experimental results

A selection of results obtained with MCPTAM deployed on the Tryphon are reported from a series of experiments conducted in the Black Box facility at Concordia University in Montréal, Canada. For the results presented in this section, the Tryphon blimp was equipped with two cameras, affixed as shown in Fig. 2. This environment is challenging due to low lighting and few stable visual features, but MCPTAM was nonetheless able to construct a map using a hand-held camera rig with two cameras. The map was post-processed using the MAPEDITOR to align the floor with the xy -plane and one of the walls with the yz -plane, both to aid in visualization and to ease the task of controlling the airship by defining an intuitive world coordinate frame. The resulting pre-built map is illustrated in Fig. 13 where the labeled locations correspond to those in Fig. 12, and are shown to aid in visualizing the results.

The following camera configurations are processed with MCPTAM to obtain the pose of the Tryphon blimp during a sample flight: camera #1 only, camera #2 only, and both cameras together but each in freerun mode (unsynchronized). The last configuration is made possible by MCPTAM's support for asynchronously triggered camera groups. Although the asynchronous triggering yields a higher effective frame rate as compared to synchronous triggering, the frame rate is still not sufficiently high for closed-loop control of the blimp. At the same time, the noise level in pose estimation is amplified when operating cameras are asynchronous. Nevertheless, the two-camera results are included here to showcase the multi-camera tracking capability.

The trajectories of the Tryphon blimp as obtained from the aforementioned camera configurations are overlayed on the pre-built map in Fig. 13, where the unsynchronized two-camera results are labeled 2U. A more insightful view of the trajectories is displayed in Fig. 14 by plotting two coordinates of the blimp position as a function of time. Several interesting observations can be made based on these figures: starting with the most obvious, the problematic performance of camera #1 which was not able to localize itself relative to the pre-built map until ≈ 20 s and even then, tracking is immediately lost until 40 s, from which point stable tracking ensues. Camera #2 generates the best trajectory throughout: it did not suffer from the unfortunate fate of camera #1 as it observed a different part of the map which happened to be easier to localize in at the start. Since 2U utilizes both cameras, it took advantage of the rapid localization of camera #2 to immediately generate a valid trajectory, although the performance of 2U configuration was at times deleteriously affected by the poor performance of camera #1.




Fig. 13. Points of the pre-built map along with the trajectories of the three tests shown in (a) top view, and (b) front view, where 2U are the results with 2 unsynchronized camera.

To validate the performance of the MCPTAM tracker with one camera, it is compared to the Tryphon position estimates obtained with the SICK laser scanner. The laser was mounted near one of the walls of the experimental space, approximately 2 m above the floor. It is noted that the accuracy of MCPTAM has been investigated in previous work, in standard laboratory conditions, with initial experiments reported in [26] and a more in-depth

Table 3

RMS of the distance error between the external SICK laser and the MCPTAM system.

Circular trajectory				
Test index	1	2	3	
RMS (m)	0.22	0.13	0.14	
Square trajectory				
Test index	1	2	3	4
RMS (m)	0.14	0.18	0.11	0.10

experimental study presented in [27]. In the latter work, sub-centimeter position accuracies and sub-centiradian angular accuracies were reported against the ground-truth measurements obtained with Vicon motion capture system. Given a completely different operating environment of the Tryphon compared to the lab environment of the previous investigations, coupled with the fact that the Tryphon is tracked relative to a pre-built map with one camera only, it is informative to once again evaluate the accuracy achievable with MCPTAM.

To this end, the SICK laser measurements were processed live to determine the position of the center of the Tryphon blimp in the xy -plane. The algorithm provides reliable estimation of the position of the centroid as long as the Tryphon blimp undergoes negligible rotational motion (pitch and roll) during the maneuver and if it stays within the beam of the laser. Calibration of both systems to reach a common reference frame was achieved by employing approximately 10% of the measurements from a dataset for each maneuver and using the standard least squares technique to estimate the relative position in the xy -plane and rotation about z -axis between the two frames. The two operations – processing of raw SICK data and interframe calibration – introduce significant error in the SICK position data. Results comparing SICK vs. MCP-TAM position data are shown in Fig. 11 for a circular trajectory flown by the Tryphon under closed-loop control (to be discussed in Section 5). The interval of flat SICK data corresponds to a period when the Tryphon was outside of the laser's range. Table 3 summarizes the RMS values of the distance error for the total of seven trajectory tracking experiments: three circular maneuvers and four square maneuvers. Given the error sources in both SICK and MCPTAM data, these results demonstrate good consistency between the experiments and reasonable agreement between SICK and MCPTAM positions.

5. Control

As the Aerostabiles project evolved and expanded over the past decade, the artistic needs for Tryphons to respond robustly imposed the requirements for accurate control in variable environments where Tryphon performances and installations are held. As




Fig. 14. Plots of Tryphon trajectories over time. Note how the camera #1 configuration (cam1) was not able to provide localization until 40 s.




Fig. 15. Flow chart of ROS architecture.

well, with the number of researchers involved with the platform on the rise, modular software development became paramount. As such, the Robotic Operating System (ROS) was adopted a few years ago to allow for easy expansion and integration with new and existing hardware and software components.

5.1. ROS architecture

The Robot Operating System is an open-source set of tools and libraries facilitating the integration of supported hardware and common software features into robotic platforms.

Tryphon's modular ROS architecture allows the majority of nodes to be executed on any station, as illustrated in Fig. 15. Each node subscribes to its input messages (I) and publishes output messages (O). The Tryphons share the same software repository and thus have access to the same nodes, providing simplified and synchronized development across the fleet. Since multiple aerobots may be flying together, an IP address parameter (target_ip) allows multiple copies of the same node to run simultaneously and to target the right individual with control messages. The bottom left node in Fig. 15 corresponds to the relative localization system *You see Me, I see You* described in Section 4.1, which requires two active Tryphons, while all other nodes are attributed to a single aerobot (and can be instantiated).

The SENSORS_THRUSTERS, FORCE_DISTRIBUTION and CAMERA_NODE must be run on-board because they directly communicate with the Tryphon hardware: sonars, IMU, cameras and motors. The right-most, red-bordered frame shows the current nodes making use of the external hardware and these must be executed on a ground station connected to the respective device. All other nodes can be executed either on-board or on a ground station.

The SENSORS_THRUSTERS node uses the I²C bus for two-way communication with the sonars, IMU, compass, propellers, batteries and electromagnet. The CAMERA_NODES use the USB camera drivers from PointGrey to convert the image data into a ROS compatible message. Since the computational cost of sending this large message increases as more nodes subscribe to it, a CAMERA_RELAY node is executed on the ground station which acts like a hub and forwards the image message. Most residencies include demonstrations or shows, and the CAMERA_PROJO node to project the camera feed to an external screen turned out to be mandatory to get a sense

of how the Tryphon sees the world. The STATE_ESTIMATOR node is set to accept pose data from either MCPTAM, the SICK laser, or an array of sonars and to fuse these with the IMU information and the compass measurements. Apart from selecting the sensing mode, its main function is to compute the state through a Kalman filter and send its estimate to the CONTROL_NODE. The latter also needs a command, or a desired state input, which is either computed by the TRAJECTORY node or defined by the user or artist inputs. The external device nodes allow human interaction either from a PS3 controller, a computer keyboard, a mobile phone's orientation (as measured by its accelerometer) or a microphone (see Section 6.2). As the project is meant to be accessible to a non-expert user, likely lacking programming and Linux OS skills, a web interface allows a user from any device connected to the local network to send input messages to the CONTROL_NODE and also to see the nodes output, such as the video stream from on-board cameras.

Finally the CONTROL_NODE generates a wrench message to the FORCE_DISTRIBUTION node (see Section 5.2.3), where command to each propeller is computed according to its location on the Tryphon and safety limits, and then sent to the SENSORS_THRUSTERS node.

5.2. Controllers

A robust and accurate control is necessary in order to plan reproducible artistic performances with Tryphon blimps or for repeatable human/robot interaction experiments. Accurate control of Tryphons is also needed for docking maneuvers which will ultimately be used to create aerial architectures by assembling multiple blimps in the air.

To date, controller development has focused on controlling a single Tryphon blimp for two types of tasks: (1) set-point regulation where the objective is to maintain a Tryphon hovering in a particular location, as for example, in the Paradoxical Sleep performance [5]; (2) trajectory tracking so that Tryphon can follow a particular path, at desired speeds, as for instance when Tryphon is commanded to do so in response to human gestures or voice commands.

Towards the above objectives, the model presented previously in Section 3 was used to design and evaluate two control laws: a discrete formulation of a PID controller and a time-domain computed torque controller. An infinite horizon linear quadratic regulator (LQR) design is also available but will not be presented here.

Table 4
PID controller gains.

Input	x	y	z	ϕ	θ	ψ
K_p	3.2	3.2	3.2	0.2	0.2	3
K_i	0	0	0.1	0	0	0
K_d	20.8	20.8	20.8	1	1	15
N	100	100	100	100	100	100

The two controllers are reviewed in the following subsections and in Sections 5.2.4 and 5.2.5, they are compared in simulation and experimentally, respectively, for the aforementioned two types of tasks. It is noted that the full state feedback needed by the computed torque controller is delivered by the Extended Kalman filter which fuses the MCPTAM pose (Section 4.2) with the acceleration and angular velocity measurements from the IMU on-board. The design of the filter will not be detailed here, but is available in the work of St-Onge et al. [21]. The output command generated by all controllers corresponds to the control forces and moments, $\mathbf{F}_C$ and $\mathbf{M}_C$ which drive the dynamics of the Tryphon as per Eq. (2). A force distribution algorithm – a necessary component of the complete controller – is briefly described in Section 5.2.3 and was implemented to convert the control force/moment outputs into commands to the individual propellers.

5.2.1. PID

Based on the decoupling of the states that results when linearizing the dynamics model of Section 3, the PID controller was designed as a combination of six single-input/single-output discrete PID control laws with respective input/output pairs: (x, F_{cx}) , (y, F_{cy}) , (z, F_{cz}) , (ϕ, M_{cx}) , (θ, M_{cy}) and (ψ, M_{cz}) . In this single-input/single-output framework, we employ the discrete formulation of the controller, obtained using Tustin method [29] with a sampling rate $T = 10$ Hz, stated in the Z-domain as:

$$\frac{U(z)}{E(z)} = K_p + K_i \frac{1}{\left(\frac{2}{T} \frac{z-1}{z+1}\right)} + K_d \frac{\left(\frac{2}{T} \frac{z-1}{z+1}\right)}{\frac{K_d \left(\frac{2}{T} \frac{z-1}{z+1}\right)}{N} + 1} \quad (16)$$

where $U(z)$ is the control output, $E(z)$ is the error input, N is first-order derivative filter divisor and, K_p , K_i and K_d are, respectively, the proportional, integral and derivative gains of the PID. The gains were tuned using the *Gazebo* simulation environment (described in Section 3.2) and the final gains on the Tryphon are summarized in Table 4.

5.2.2. Computed torque controller

During performances, the Tryphon is often asked to follow scripted maneuvers, thus requiring a controller with trajectory tracking capabilities. Therefore, a computed torque (CT) controller was designed since it is theoretically able to attain zero steady state error for any trajectory achievable by the aerobot. This controller makes use of the full nonlinear model to compute the control vector τ_C as a function of the desired pose $\mathbf{q}_d = [\mathbf{p}_d^T, \boldsymbol{\phi}_d^T]^T$, velocity $\dot{\mathbf{q}}_d$ and acceleration $\ddot{\mathbf{q}}_d$, leading to the following control law [30]:

$$\tau_C = \mathbf{M}_q (\ddot{\mathbf{q}}_d + \mathbf{K}_v (\dot{\mathbf{q}}_d - \dot{\mathbf{q}}) + \mathbf{K}_p (\mathbf{q}_d - \mathbf{q})) + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \dot{\mathbf{q}} + \mathbf{G}(\mathbf{q}) \quad (17)$$

where (with $\mathbf{S}$ from Eq. (1))

$$\mathbf{M}_q = \begin{bmatrix} (m + A_m) \mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & (\mathbf{J} + \mathbf{A}_j) \mathbf{S} \end{bmatrix} \quad (18)$$

$$\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \dot{\mathbf{q}} = -\tau_l - \tau_v + \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ (\mathbf{J} + \mathbf{A}_j) \dot{\mathbf{S}} \end{bmatrix} \dot{\mathbf{q}} \quad (19)$$

$$\mathbf{G}(\mathbf{q}) = \begin{bmatrix} 0 & 0 & 0.012 \int_0^t (z_d - z) dt & 0 & 0 & 0 \end{bmatrix}^T \quad (20)$$

and $\mathbf{K}_v$ and $\mathbf{K}_p$ are symmetric positive definite gain matrices designed to compensate for measurement and model errors. Again, after some tuning using the simulator (see Section 3.2), these are defined as:

$$\mathbf{K}_v = \begin{bmatrix} 0.42 \mathbf{I}_3 & \mathbf{0}_3 \\ \mathbf{0}_3 & 0.49 \mathbf{I}_3 \end{bmatrix}, \quad \mathbf{K}_p = 0.6 \mathbf{I}_6. \quad (21)$$

In a standard formulation of the computed torque [30], the vector $\mathbf{G}(\mathbf{q})$ is defined to compensate for the known force and moment created by gravity. However, in the dynamics of Tryphon, the imbalance force $\alpha \mathbf{F}_g$ is unknown and changes when helium or weights are added to balance the Tryphon. On the other hand, the moment created by buoyancy, although could be compensated, provides a stabilizing effect on the pitch and roll dynamics of the blimp. For these reasons, the “gravity” compensation term $\mathbf{G}(\mathbf{q})$ is defined as per Eq. (20), with a single integral term in the z-direction to eliminate the steady-state error in height.

5.2.3. Force distribution algorithm

Since the Tryphon is either equipped with eight or twelve propellers, while a rigid body free to move in 3D space has six degrees-of-freedom, the force distribution problem is an under-determined linear system and thus, does not have a unique solution. To resolve this indeterminacy, the force distribution is conceptualized by first evenly distributing the command force components in x-, y-, z-directions between the two x-, two y- and four z-acting propellers, respectively. Then, the additional couples to generate the desired command moment $\mathbf{M}_C$ are computed and added to the propeller thrusts. The resulting force distribution solution is thus computed using the following matrix equation:

$$\mathbf{F}_p = \frac{1}{4} \begin{bmatrix} 2 & 0 & 0 & 0 & 0 & -\frac{1}{l_x} \\ 2 & 0 & 0 & 0 & 0 & \frac{1}{l_x} \\ 0 & 2 & 0 & 0 & 0 & \frac{1}{l_y} \\ 0 & 2 & 0 & 0 & 0 & -\frac{1}{l_y} \\ \frac{l_z}{l_x} & -\frac{l_z}{l_y} & 1 & -\frac{1}{l_x} & -\frac{1}{l_y} & 0 \\ \frac{l_z}{l_x} & \frac{l_z}{l_y} & 1 & -\frac{1}{l_x} & \frac{1}{l_y} & 0 \\ -\frac{l_z}{l_x} & \frac{l_z}{l_y} & 1 & \frac{1}{l_x} & \frac{1}{l_y} & 0 \\ -\frac{l_z}{l_x} & -\frac{l_z}{l_y} & 1 & \frac{1}{l_x} & -\frac{1}{l_y} & 0 \end{bmatrix} \times \begin{bmatrix} \mathbf{R} & \mathbf{0}_3 \\ \mathbf{0}_3 & \mathbf{I}_3 \end{bmatrix} \tau_C, \quad (22)$$

with the length parameters l_x , l_y and l_z defined in Section 3.1.2. Since the propellers' commands sent to their ESCs are in percentage of the maximum velocity, steady-state characterization tests were conducted to determine the force-velocity relations of the propellers (see Section 2.1.2). Thus, the actual propeller commands are computed by inverting these input-output relationships for all components of $\mathbf{F}_p$. There is also a minimum and maximum saturation function which scales the magnitudes of all thrusts, while preserving the direction of the desired control wrench. Furthermore, a filter is implemented which limits the maximum rate of command increase per propeller, because very large jumps in motor commands tend to crash the system.

Table 5
Simulation results for trajectory tracking control (perfect feedback).

	PID	CT
RMS distance error (m)	0.0498	0.0886
MAD roll error (rad)	0.0084	0.0079
MAD pitch error (rad)	0.0087	0.0074
RMS yaw error (rad)	0.0005	0.0004
Δz (m)	0.0359	0.0499




Fig. 16. Circular trajectory with PID and computed torque tracking in simulation (perfect feedback).

5.2.4. Simulation results

The performance of each controller in our *Gazebo* simulation is now addressed. We compare the PID and computed torque (CT) controllers on a trajectory tracking task for the reference trajectory: a 2.5 m radius circle (see Fig. 16) traveled at a constant velocity of 0.075 m/s. To quantitatively compare the performance of the two controllers, the five selected metrics are: the root-mean-square (RMS) error in the plane of the circle $\sqrt{(x_d - x)^2 + (y_d - y)^2}$, the RMS error in yaw, the mean absolute deviation (MAD) in the roll and pitch angles of the blimp and the difference between the maximum and minimum z positions (Δz), all of these are summarized in Table 5.

The two controllers show similar performance for all metrics, although the PID outperforms the CT for translational tracking, whereas the CT has a slight advantage with orientation. The PID maintains its advantage in position tracking performance even when the actuator constraints are removed in simulation. We attribute this somewhat surprising result to the better tuning of the PID control gains, coupled with the rather slow (i.e., not very dynamic) motion of the blimp.

5.2.5. Experimental results

The performance of each controller in experiments is now addressed. We begin by comparing the PID and CT controllers for set-point control with a square trajectory. The four set-points are the four corners of a 5-meter side square, defined parallel to the xy -plane, at a height of 2.5 m (see Fig. 17(a)). The controller switching between the set-points occurs when the distance between the Tryphon's actual position and the desired current corner position is less than 0.20 m. In this scenario, however, the performance at subsequent corners of the square is affected by the accuracy of regulation achieved at the previous set-point (corner). To compare the performance of the two controllers, the six selected metrics are: the overshoots at the first three set-points in the direction of motion (e.g., the first overshoot is the maximum error in the X -direction), the RMS yaw error, Δz , the sum of the propeller's commands and the time to complete the square. The last metric

Table 6
Results for set point control.

	PID	CT
Overshoot at 2 (m)	1.26	1.11
Overshoot at 3 (m)	0.23	0.25
Overshoot at 4 (m)	0.04	0.1
RMS yaw error (rad)	0.03	0.06
Δz (m)	1.78	1.18
Sum commands (N)	70154	69717
Time (s)	134	136

Table 7
Results for trajectory tracking control.

	PID	CT
RMS distance error (m)	0.070	0.117
MAD roll error (rad)	0.0119	0.0108
MAD pitch error (rad)	0.0120	0.0093
RMS yaw error (rad)	0.011	0.22
Δz (m)	0.42	0.42

is representative of the settling time of the two controllers. Note that the RMS yaw error and the max-min height difference are computed over the complete maneuver and are included here to give a more complete picture of the controllers performance. The “square” experiment was executed three times with both controllers and the average results (averaged over the three experiments) for the metrics are presented in Table 6. The results show comparable performance for the two controllers, with the exception of the first overshoot and the Δz metrics for which the CT controller performs better. Considering the CT controller provides nominal cancellation of the drag on the blimp, it is reasonable that the overshoot at the first way point (corner 2) with CT is less than that with the PID controller.

A second set of experiments was carried out to compare the PID and the CT controllers on the same trajectory tracking task as used in simulation in Section 5.2.4: a 2.5 m radius circle (see Fig. 17(b)) traveled at a constant velocity of 0.075 m/s. The performance metrics remain the same as for the simulation. The circle was executed three times with both controllers and the average values are presented in Table 7. The two controllers show similar performance for all metrics except for the RMS distance error where the PID outperforms the computed torque, similarly to the results of the *Gazebo* simulation. The experiments also show superior yaw tracking with the PID. Overall, these experiments highlight the main source of difficulty in controlling the Tryphon blimp which is its high sensitivity to variations in ambient conditions. Operating the blimp in large spaces comes with the challenges of dealing with local disturbances on the blimp, as for example, local air drafts. Neither the PID nor CT controllers are robust to these disturbances, which can lead to lack of repeatability between experiments and unexpected deviations from the desired motions. Example of the latter is the noticeable drop in the height of Tryphon visible in the region $(x, y) = (2, 0)$ for both the square and circular trajectories.

For the purpose of presentations of Tryphon to an audience, as in the experiments on Emotions of Movement to be presented in Section 6, the CT controller was chosen because of its superior stabilization of the roll and pitch oscillations; these are more noticeable to an audience than errors in translational tracking.

6. Human-robot interaction

The focus of this research is to enable artists and performers create artwork with a reliable and flexible flying robotic platform. Results of Section 3 allow for a better understanding of the dynamics of our uniquely shaped blimps through accurate simulation and visualization. Section 4 detailed the techniques for Tryphons to localize relative to each other and to their environment, in




Fig. 17. (a) Set point regulation with PID and CT controllers, (b) Circular trajectory with PID and computed torque tracking.

various residency setups, and thus led to their autonomy. Finally, Section 5 introduced two approaches to behavior and control coding that allow for fast integration of new paths, reactive response to perturbations and controller design for different performance contexts. It is only with the potential and robustness enabled by these technological developments that reliable interactions with artists can be designed.

Since the interactions with Tryphons are explored within the artistic realms, an unusual research paradigm emerged based on the developments of *emotions*-driven interactions, rather than remote control approaches. Standard robot interaction devices such as gamepads or cellphones do not present much interest for artistic performances. Indeed, the relation between the aerobots and the humans must be intrinsically organic, and not only perceived as such. For this reason, a major contribution of this research lies in exploring different devices to enable human–Tryphon interactions with the ultimate goal of developing original modes of interaction that can be used naturally by non-experts during actual artistic performances. In the following subsections, four such systems are outlined. Before introducing these means of interpreting the human emotions and intentions, the potential of expressiveness of the Tryphons for emotions is first addressed.

6.1. Emotions of movements

As previously stated in Section 2.2, artists explore the movements of the Tryphon by manipulation of the structure and, through their aesthetic inputs, many public shows have been created over the course of the Aerostables project. Amongst the performances and installations presented in the past several years, feedback from the audiences frequently pointed to the perceived intention of the aerobot [5]. Since no such intentions were programmed, nor do Tryphons have anthropomorphic morphology which could inspire intentions, these perceptions were caused by their movement. Thus, as is the case for most robotic platforms, trajectory design was found to be of critical importance for the interaction of the Tryphons with humans [31]. Unique to our platform, however, is the use of these trajectories to impart emotional expression to the Tryphons. The importance of perceived emotions is well-known in HRI [32,33], but it is usually explored in the context of animal-like or human-like robots.

Helped by the choreographers and visual artists participating in the project, a first vocabulary of intended feelings and expressions was created from sequences of displacements and rotations, as well as with the precise dynamics (acceleration and velocity) of these sequences. The underlying assumption was that *emotions* can




Fig. 18. 4 sequences studied for their emotional interpretation potential as designed from the insight of a choreographer and a media artist. (a) is viewed from the side while all the others are viewed from the top. A schematic person is at the right of each sequence to give a sense of distances traveled by the aerobot.

be inferred by the audience through perception of the *movement* only of a non-anthropomorphic robot.

The study was conducted in the Black Box facility of Concordia University, the location of earlier experiments on MCP-TAM. It is a large room with a black floor, three black walls and one white wall used for projection, as well as stage lighting fixtures suspended from the ceiling (see Fig. 12 for an overview of the flight area). Based on the hypothesis that had to be verified, four motion sequences were proposed, as shown in Fig. 18:

- In the first one, the Tryphon slowly moves backwards, starting from the center of the room; when it gets close to the back wall, it flies up towards the ceiling moving slightly faster. After a short pause, it dives directly and rapidly straight towards the audience, following a diagonal line. This sequence was meant to be aggressive, as if the Tryphon was attacking people in the audience.
- In the second sequence, the Tryphon hovers above the audience, oscillating in yaw from right to left with a period of about one minute. In this sequence the aerobot was meant to be curious, as if it was scanning and analyzing the audience.
- In the third sequence, the Tryphon starts at mid-height of the room, then heads towards the front right corner of




Fig. 19. (a) Stressed versus Relaxed for the 4 sequences studied, (b) Attacked versus Protected for the 4 sequences studied, (c) General perception of the Tryphon over all sequences in terms of its anthropomorphism. The crosses are statistical outliers. The vertical axis is the score following the extrema of the x-axis labels.

the room, thus getting closer to the audience, then flies towards the opposite corner while rotating through 180°. This sequence was meant to represent fear, as if the robot was approaching a new danger and then quickly fleeing.

- The fourth sequence, identical to that used in Sections 5.2.4 and 5.2.5 shows the Tryphon slowly describing a large circle at a low altitude. This last sequence was meant to be calm, as if the Tryphon was tamed and comfortable in the presence of the audience.

The interpretation of emotions or intentions through displacements and rotations is very difficult to quantify. For instance, it strongly depends on the cultural background of the audience [34]. Such cultural biases are impossible to avoid, but their impact can be evaluated and taken into account by selecting a sample of observers that is representative of a precise social group.

The four movement sequences were sequentially shown to an audience of 20 persons (10 male, 10 female). The audience was mostly composed of students, about a half with arts or humanities background, and the other half from engineering or sciences. They were aged between 17–50 (mean = 28, stddev = 7) and 13 spoke French while the other 7 spoke English. These observers were handed a questionnaire in which they were asked about their interpretation of the different sequences, focusing on the anthropomorphic aspects of the Tryphon's cubic morphology and its behaviors. Participants were asked to evaluate for 5 different axes inspired by the work in [35], on a scale from 1 to 10, to which extrema their perception was closer. The results of the group's responses are graphically represented with the box-and-whisker plots in Figs. 19–20. Before interpreting the results in terms of the average and variability of the scores of the sequences perception by the participants, the statistical effect of the sequence on the mean value of each interpretation axis must be addressed. Since the sample is relatively small (<30), a non-parametric study was selected, the Friedman's test. Using its implementation in the Statistical toolbox of Matlab, the p -value of the sequences effect on the mean of each interpretation axis is 4.15×10^{-4} , proving that the influence is statistically significant (<0.001).

As shown in Fig. 19(a), the Tryphon was perceived more calm and more quiescent in the last sequence. This result was the one expected, especially since the noise level of the motors during this slow displacement was very low and regular. Fig. 19(b) shows that the audience felt more protected (least threatened) during the three last sequences than during the first one. Again, this confirmed our expectations on the Tryphon threatening the audience. As for the second and third trajectories, the expected perceptions were not confirmed. For instance, the curious state of the second trajectory left the audience feeling less observed than with the first and last ones as shown in Fig. 19(c). This may be explained by Tryphon's height during these movements creating a detachment of the audience. Fig. 20(c) reveals that the audience's average perception of the Tryphon is not anthropomorphic, nor is it considered

as completely 'fake' or 'machinelike', but rather in between. This appears in contradiction with the qualitative results gathered during previous installations, where anthropomorphic interpretations were frequent [5]. However, differently from the setup used in the present study, those installations were interactive, which tends to increase the will to develop relations with artificial objects. The lack of direct interaction in this particular setup was deliberately chosen in order to focus strictly on a trajectory/emotions lexicon. Thus, by simply following scripted trajectories the Tryphon is not framed into a machine-like only perception. While most participants commented on the movements of the Tryphon in terms of trajectory in the space, some, even in this non-interactive context, described its behavior as if the robot was 'observing' them or being 'shy'. The audience unanimously described the movements of the *Aerostabile* as 'elegant', which is important since this criterion is one of the artistic design objectives from the outset of the *Aerostabiles* project.

These results clearly demonstrate that triggering the perception of a specific intention or emotion using only the movement of an abstract geometry is possible. Such a postulate may be an important paradigm shift in designing interactive robots. Supporting the guidelines drafted in [36], instead of focusing on the robot's geometric design and appearance, the designers may consider placing more emphasis on its dynamics.

6.2. Song triplets

The interpretation of the cubes' moods by visitors or by an audience is, as mentioned earlier, related to cultural background, but also to the moods of the visitors themselves. This leads to a situation where the spectrum of possible interpretations can be rather wide, and the final interpretation unpredictable because of the combination of individual and socio-cultural factors. Unpredictability does not forbid investigation, however, and one of the ways chosen to explore the impact of this "cultural dialogue effect" is through human voice. A software module that allows the control of the cubes by human voice, through short 3-note melodies sung by a performer, or by anyone with minimal singing skills was implemented. The expressive potential of the human voice, combined with the general mood of each of these melodies (major, minor, attack, envelop, etc.) installs an initial atmosphere in which the reactions of the cubes take on different interpretations than in full silence.

The melodies could be sung starting at any note: the information considered by the analyzing software are not the notes themselves, but rather the intervals between the notes, in order for people not endowed with absolute hearing (the vast majority of people) to use the system. The use of note triplets extracted from familiar songs or nursery rhymes facilitates the initial training of the performers. Subsequently, triplets can be specifically composed for the performances, in order to match the general ambiance sought by the artistic team.




Fig. 20. General perception of the Tryphon over all sequence in terms of its anthropomorphism. The crosses are statistical outliers. The horizontal axis is the score following the extremum of the y-axis labels.

Table 8

Success rate for 6 sung triplets. For each triplets, the series of 3 notes are indicated on the horizontal axis according to the latin nomenclature. The numbers following the notes are their octaves.

Triplet	Success rate (%)
LA3, SI3, DO#4	40
DO#4, SI3, LA3	60
LA3, DO#4, DO4	20
LA4, SI4, DO#5	100
DO#5, SI4, LA4	100
LA4, DO#5, DO5	80

The latest version of the triplet extraction software is integrated in a ROS node and is based on the aubio¹⁴ library for Linux. The basic principle of the software is to detect the different notes of the short song. This is done by analyzing the pitch of each note, but also the variations revealing discrete sound events, which allowed to individualize them. The variations are detected by evaluating the local energy differences between windows in the signal. If the measured energy level is under a threshold, usually low, but necessary since most singers do not make a clear stop between notes, the system then considers the pitch which is detected based on the algorithm from the work of Brossier [37]. To minimize false detection, a timeout triggers the reset of the triplet detection if three notes are not detected on time. This specific parameter has to be adapted to the pace of each performer.

The results shown in Table 8 are obtained from an experiment that was conducted with 10 participants, using the same parameters of the extractor each time. A trial was considered a success if the analyzing software recognized the triplet and output the correct command. It is reasonable to state that the software would achieve better performance if its parameters were finely tuned to match each individual's vocal characteristics. Nevertheless, these results show great performance at the 4th octave, even for a more difficult set of intervals (LA, DO#, DO). Poor performance starts to appear at the 3rd octave (shown by the first three sequences in Table 8) and was empirically found to worsen below that range of frequencies. As a matter of fact, most note detection algorithms become unstable as the pitch decreases, since the differences between notes in terms of frequency diminish quickly. With synthetic notes, the algorithm has a success rate of 100% on a wider spectrum. Consequently, part of the challenge was to adapt to different voice gains and specific texture as well as to deal with continuous melody without clear stops between notes and sharp note attack. The system was inaugurated in Sao Paulo FILE festival in 2012, where a performer was controlling the Tryphon movement through a set of 6 different triplets associated with scripted behaviors. After hours of fine tuning the system parameters to the performer and with the help of a professional noise reducing microphone, 10 shows were staged with the same set of commands without failures.

6.3. Distance and light sensors

Starting with their predecessors Mascarillons (2005–2010), the Tryphons have been equipped with various sonars to detect obstacles and walls at close distances (<6 m). A PI controller maintains the distance from a wall stable and is overwritten by an extreme reaction control (full thrust) when an object is detected too close to the blimp (<0.75 m). In this way, the sonars created around the *Aerostabile* robot an area that the performers came to perceive and identify as an “intimacy sphere”: an area inside which the Tryphon appears afraid, reacts in order to avoid potential collisions and keeps a comfortable distance from anything surrounding it. The sonars allowed two or more performers to easily control the Tryphon translations by performing in the area equivalent to a 12-meter diameter circle.

The sonar sensors are also equipped with light intensity detectors. The light detection ability can be used in a similar manner, however, it provides more reactivity to the environment manipulation (projectors, sun-light through windows and ceiling lamps) and thus, is harder to adjust. A 2009 installation in Moscow, called “Geometric Butterflies”, was precisely based on this feature. It involved three aerobots hovering in a flight area surrounded by dark blue, intense spotlights. Drifting close to the lights, the *Aerostabiles* were pushed back towards the center of the area, where they got close to each other; their “intimacy sphere” algorithms then repulsed them back towards the spotlights, and so on, in an endless dance which they described during three weeks of never-repeating orbits and long-lasting trajectories. With these interfaces, however, global positioning of the Tryphons is not known and therefore, complex emotion-driven trajectories are not possible. The motion of the aerobots is thus managed with the sonar beams acting analogously to puppets actuated with strings.

6.4. Myo armband

In order to create a two-way emotion-driven interaction, the emotions of the performer need to be recognized. Most solutions for such detection use close range extraction of facial expressions from a camera or external motion capture system, neither of which are compatible with this project. Since 2014, an accessible and powerful gesture control armband, called the Myo,¹⁵ has been available on the market. The Myo includes a 9 degree-of-freedom IMU as well as 8 electromyogram sensors (EMG) disposed symmetrically around the forearm. The set designers and choreographers working on the project quickly saw the potential of this device. In a first iteration to explore its usability, it was interfaced with a series of robotized spotlights so that the beams could follow the movements of the dancer, and match their dynamics with color and intensity changes. The spot orientation was a direct mapping

¹⁴ <http://aubio.org/>.

¹⁵ <https://www.thalmic.com/>.

of the filtered orientation information from the IMU only. The color was set to evolve in a bright yellow to dark green range following the sum of the energy measurements from the EMGs.

The success of this qualitative exploration, confirmed by the ability of the performers to create discernible patterns with the spot orientation and color, lead to the development of a sophisticated and automated EMG recognition system. The goal was to achieve a robust and credible feeling of interaction between the *Aerostabile* robot and the performer, by mapping complex choreographic sequences to specific EMG patterns. A first lexicon of movements was determined during the last residency, with the help of a choreographer and a dancer, based on the dancer's movements influence on the lighting (as describe above). Many similar datasets were created, from which different features were engineered, in order to train a classifier with the most relevant information available.

While the data flow from the IMU is easy to handle and interpret (as it measures physical quantities associated with movement), the EMG data are noisy and impossible to directly associate with a specific movement. Based on the general similarity between the accelerometer data from a surface inspection tip device in [38] and the EMG signals of the Myo, an initial set of features extracted from the device was selected for study. It consisted in the following:

- minimum and maximum values,
- mean, variance, skewness, kurtosis and the fifth moment,
- Integrated Values (IV—sum of absolute value of all EMG or orientations),
- Mean Absolute Value (MAV),
- MAV1 (weighted average with more weight given on central values),
- MAV2 (as MAV1, with a different weight distribution),
- Root Mean Square (RMS),
- Zero Crossing,
- Waveform Length (WL),
- Slope Sign Change,
- Willison Amplitude (WAMP—number of overreached difference in EMG signal amplitude),
- Median Frequency,
- Frequency-Based Features (frequency and amplitude of most powerful component of the power spectrum extracted from non-overlapping windows of 50 samples. For a time-window of 1 s (200 samples), this resulted in 4 pairs of frequency–amplitude. A Hamming window was used to avoid frequency leakage).

As detailed in [39], four classifiers were compared (Adaboost, Random Forest, SVM-I and SVM-II) and the selection of the most relevant features was conducted by comparing the results of successive classifications over a dataset with different sets of features. The resulting final features are:

- for IMU orientation: Maximum, Mean, Variance, IV, MAV and RMS;
- for EMGs: FBF, Variance, IV, MAV and RMS.

A second experiment was conducted only with dancers, to assess the performance of the classifier. A key aspect of this experiment was the ability to train this classifier online, based on a fast implementation in Python with scikit-learn. This allowed the performers to develop choreographic sentences on-the-fly using their own conceptual categorization, i.e., emotions or ambiances, with the classifier trained live to discriminate each class. To increase the richness of interaction, two Myo armbands were worn by the performers, one on their forearm and the other on their calf. The final settings for the classification performances showed in Fig. 22 were:

Actual	1	90.7	6.1	3.3
	2	1.1	98.8	0.1
	3	1.0	3.6	95.3
		1	2	3
		Predicted		

Fig. 21. Confusion matrix (in %) for RF trained on 3 moods.




Fig. 22. RF 9-class performances on the second experiment. (a) means arm and (l) means leg.

- Random Forest algorithm with the previously detailed hyper-parameters;
- 1 s time-window to detect a class;
- overlaps of 0.3 s between successive windows;
- measurements of the orientation of the right forearm and right calf, with 4 EMG on the right forearm.

It resulted in a classification rate of 94% (average on 100 runs on each performance) on the 3 moods defined by the artists for the experimental choreography. Fig. 21 shows that few blending occurred, with most of the errors arising during transitions. With this level of performance, we can conclude that our approach is suitable to closing the loop with a robot, and in some sense bestowing it the ability to detect which emotional ambiance the user is choosing to express.

The informative value of the EMG from the arm EMG_a against the IMU of both the leg IMU_l and arm IMU_a was compared by the accuracy of the RF on the 3 classes. The classification results for all possible combinations is shown in Fig. 22. The IMU, as a single source of information, performs better than the EMGs, which is not surprising given the complexity and noise of the latter. Taken separately, the IMUs on the forearm and calf give similar performances but together they increase the performance of the classifier by more than 10%. Finally, combining the two IMUs with the forearm EMGs increases the classification results by at least 5 percentage points.

The previous experiments, which were driven by the interest from the artist partners of the project, show great performance of an emotionally-based classification of body movements. It is

now integrated in a close loop with our ability to convey emotions through the Tryphon movements (Section 6.1) to create an improvised and hybridized “pas de deux”.

6.5. Curiosity module

Detection of performers intentions can also come from a global scene interpretation. Doing so does not require from the performer to always stand in the field of view of the camera but rather to deal with the attention of the aerobot. The system developed, referred to as the curiosity module, comprises of a pan-tilt camera unit controlled by an interestingness detection algorithm based on ROST [40], a realtime spatiotemporal topic modeling technique. ROST has been previously used for developing curious exploration robots that detect and gather data about interesting phenomena in the world [41,42]. The general idea behind the use of ROST for modeling curiosity in a robotic platform is the following. ROST learns a decomposition of the observations in terms of high level scene constructs (called topics), in realtime. For each low level observed visual word w , it computes a high level label z , which takes into account the spatiotemporal context. The likelihood of the observation w is modeled as

$$P(w|x, y, t) = \sum_z P(w|z)P(z|x, y, t). \quad (23)$$

Here x, y, t are the spatiotemporal coordinates of the observation w . The distributions $P(w|z)$ and $P(z|x, y, t)$ are modeled using Dirichlet priors, to encourage sparse representations. Given the incoming image data described using of this model, the curiosity module then identifies parts of the image that are most perplexing (i.e., with high uncertainty). Using a PID controller, the module then controls a pan-tilt unit to track these interesting features, and to keep them centered in the camera field of view. To avoid getting caught in a local optima, the camera is reoriented to a randomly chosen part of the scene after every few minutes.

During a recent public installation at the Chalet Mount Royal in Montreal, Canada, experiments were conducted with the use of this unsupervised machine learning technique to automatically detect and record what is interesting in the scene, and to broadcast it live. The curiosity module controlled pan-tilt camera unit was mounted on the Tryphon blimp, which was manually flown over and around the audience and the performers, producing the effect of a curiosity observer. The pan-tilt unit autonomously controlled the camera to search for the most interesting feature in the crowd below and its ‘search’ was projected on a large screen for the audience to follow its quest.

Fig. 23 shows examples of what the curiosity module found interesting. The curiosity module enables a novel form of interaction with the audience and the performers. Projecting the output of the curiosity module on a large screen at the venue blurs the line which separates the performers and the audience by presenting an impartial view of what is interesting. The artists thought already of many usages of the curiosity module apart from the interesting strangeness of projecting the live view of such a system. The loop can be closed in a multi-dancers performance to challenge the attention of the aerobot and play on the lexicon of movements it may react to as opposed to the one it sees as boring. The projection of its view also gives a new perspective on the choreography.

6.6. Experimental observations

The comprehension of the system dynamics and the potential created by the integration of multiple sensors on the Tryphons and on human performers led to various kinds of exhibitions and artistic performances. The ones appearing early in the development of the project are extensively described in [5]. A key

observation worth repeating is that randomly selected visitors attributed different psychological profiles and different intentions on identically programmed Tryphons. The movement of the robots, fighting against air streams and other physical perturbations are not exactly the same depending on where they are situated in a given room. This influences their dynamics, which is then interpreted differently. A performance involving video projections on the flat surfaces of the blimp was made possible without extensive development of a video mapping system, because of their robust and stable control as described in Section 5. The video projection was fixed and the Tryphon was perfectly stabilized into it. The visitor could then approach the hovering robot and engage in a discussion with it. This art experiment followed a Wizard-of-Oz model, with a hidden actress discussing with the visitors while mimicking the schizophrenic intelligence of a flying automaton. The projection of the actress eyes on the cubes faces made it more anthropomorphic, but it still looked uncanny as a 2 m cubic robot. Nevertheless, many visitors tried to teach it poetry or songs and even came in for long discussion, as if it was a reliable confidant.

More recent events in Montreal and in Sao Paulo made use of the knowledge acquired on the perception of the Tryphon movements, together with new interaction devices to create unique public demonstrations and performances. A performance involving 2 Tryphons and 2 dancers improvising based on a fixed set of choreographic sentences was conducted in front of a small audience (20 people) in Montreal. The performers and choreographer studied the dynamics of the aerobots from the team knowledge and infer from it a lexicon of movements. In Sao Paulo, a single performer was dancing and controlling the Tryphon by singing triplets commands. The visitors were numerous (almost 100 per day for 4 days) and mystified by the accuracy of the performer control over the slow and large hovering aerobots. Its movements were scripted from the previous observations of human perception and triggered by specific triplet whose volume set the amplitude and velocity of the trajectory. As stated earlier, the interactive setup showed great reliability and reproducibility. Another performance, this one involving four dancers and two Tryphons, was held in a summer pavilion on top of Montreal Mont-Royal. The airstream were hard to deal with and the interference created by the mobile devices of the 400 people audience created a complex context of work. Nevertheless the Tryphons behaved as uncanny animals, difficult to tame because of their fear of humans being. This was perceived by all attending people interviewed.

As briefly introduced earlier, the robustness of the system also allows for artistic performances not involving interaction with performers. For instance, knowledge of the Tryphons dynamics with their powerful propellers made possible an installation in the hall of a Quebec museum, with huge airstream from the main entrance doors and localizing with distance sensors relative to small flat features in the room. The PS3 controller together with on-board distance sensors feedback and compass controlled yaw, serve for the exploration of a 200 m wide cavern in south of France (La Verna). The on-board cameras were streaming exclusive images of unreachable rock walls by humans.

7. Conclusion

In this paper, recent work completed by the multidisciplinary research team working on the Tryphon robotic blimps, the most recent prototypes of the *Aerostabiles* project, was presented. Constantly oscillating between the dry transparency and sheer precision of the mathematical equations to the open, poetical interpretations of artificial beings deliberately confused with conscious, interacting biological organisms, the project progressively managed to develop and intensify contacts, communication and exchanges between cultural fields that are generally considered



Fig. 23. Examples of interesting observations captured by the curiosity module, mounted on Tryphon. The interesting parts of the scene are highlighted with red color, and the most interesting part of the scene is marked with the white blob overlaid on real images. On the left we see the curiosity module fixated on the performers as they do their routine, and on the right we see the module being interested in an audience's eye t-shirt. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

alien to each other. It directly addresses the problematics of the relation between strictly formal events, such as the spatial rotation of an object, and the emotional interpretations that such events can trigger. Through multiple exchanges and experimentation, the *Aerostabiles* robots opened a huge territory of investigation in which groups coming from diverse disciplines learn to cohabit, and over which the flying cubes stand as the catalysts of disciplinary and methodological exchanges, transforming their own technological identity into a full, extensively poetic medium, becoming a hub around which new modes of discussion can take place.

Since the Tryphons are meant to perform with actors, dancers, singers or with their artificial siblings in front of different kinds of audiences in artistic contexts, their development is first and foremost driven by the needs of the artistic team; this team in return nourishes itself and enriches its creative potential with the growing technological possibilities of the Tryphons. The dynamics model of the automata, constructed and empirically validated by the engineers, progressively emerges from the exploration by the artists of the complex flight dynamics of these cubic flying machines. From this model, the engineers designed two different types of controllers.

For the control to be accurate, the localization system has to be reliable while in tune with the context and constraints of real-life artistic installations: to minimize expensive or extensive setups. The *You see me, I see you* relative localization system was introduced, as well as the multi-camera parallel tracking and mapping system. The performance of the latter employed as feedback to the Tryphons controllers was explored and compared to an external SICK laser. It was shown that the PID and computed torque controllers, both have good regulation and trajectory tracking performance, but also revealed the need for more robust controller design. In the most recent public installation, the controllers showed appropriate performance and reliability for the project uses. However, the MCPTAM system was found to be vulnerable to theater spotlights and sunlight rays from windows. Finally, four interaction situations typical of the artistic context were demonstrated with experimental results and examples of performances.

While further pursuing research on robust localization systems, possibly considering the use of on-board 3D scanners in the future, investigations of controllers now focus on the study of rendezvous trajectories for automatic docking of two or more hovering Tryphons. In terms of interactions, new technological devices are investigated. More importantly, future performances are being planned to fully integrate all the HRI potential of the *Aerostabiles*.

Acknowledgments

The authors would like to acknowledge the financial support of the Quebec Ministry of Culture (MCCCF), the Quebec Council for Arts and Letters (CALQ), the Canadian Council for the Arts, the Fonds Québécois de Recherche sur la Société et la Culture (FQRSC), the Fonds Québécois de Recherche sur la Nature et les Technologies (FQRNT) (165565), the Hexagram Institute, the Interuniversity Centre for media Arts (CIAM), as well as of the UQAM fund for support of research and creation (PAFARC).

References

- [1] F. Popper, *Art of the Electronic Age*, Thames and Hudson, New York, US, 1993.
- [2] E. Kac, *Foundation and development of robotic art*, *Art J.* 56–3 (1997) 60–67.
- [3] D. St-Onge, C. Gosselin, N. Reeves, [voiles—sails]: a modular architecture for a fast parallel development in an international multidisciplinary project, in: *Proceedings of IEEE ICAR 2011*, Shanghai, China, 2011, pp. 482–488.
- [4] C. Cooke, See in particular Krutikov's flying cities, in: *Editions of The Museum of Modern Art*, New York, USA, 1990.
- [5] D. St-Onge, N. Reeves, *Human Interaction with flying cubic automata*, in: *Proc. of the 5th Conference on Human–Robot Interaction, Workshop: "What do collaborations with the arts have to say about HRI?"*, IEEE/ACM, 2010.
- [6] D. St-Onge, N. Reeves, *Sails/voiles project website*, 2010. <http://www.voiles-sails.org> (Accessed 30 September 2014).
- [7] A.d.I.E.L. Moreno, J.M. Armingol, S. Garrido, M.A. Salichs, A genetic algorithm for mobile robot localization using ultrasonic sensors, *J. Intell. Robot. Syst.* 34 (2002) 135–154.
- [8] L.R. Müller, A. Rottmann, W. Burgard, A probabilistic sonar sensor model for robust localization of a small-size blimp in indoor environments using a particle filter, in: *Proc. of the IEEE International Conference on Robotics & Automation, ICRA*, Kobe Japan, 2009, pp. 3589–3594.
- [9] R.H. Bishop, *The Mechatronics Handbook*, 2 Volume Set, second ed., CRC Press, 2002.
- [10] Y. Li, M. Nahon, Modeling and simulation of airship dynamics, *J. Guid. Control Dyn.* 30 (6) (2007) 1691–1700.
- [11] A.I. Korotkin, *Added Masses of Ship Structures*, Springer Science & Business Media, 2008.
- [12] C. Georgiades, *Simulation and control of an underwater hexapod robot* (Master's thesis), 2005.
- [13] S.F. Hoerner, *Fluid-dynamic drag: Practical information on aerodynamic drag and hydrodynamic resistance*, in: *Hoerner Fluid Dynamics*, 1965.
- [14] H. Ghassemi, E. Yari, The added mass coefficient computation of sphere, ellipsoid and marine propellers using boundary element method, *Polish Maritime Research* 18 (1) (2011) 17–26.
- [15] H. Mazhar, Experimental investigation of added masses of parallelepipeds on shallow water.
- [16] J.A. Deruntz, T.L. Geers, Added mass computation by the boundary integral method, *Internat. J. Numer. Methods Engrg.* 12 (3) (1978) 531–550. <http://dx.doi.org/10.1002/nme.1620120312>.

- [17] G. Birkhoff, A. Schoenstadt, *Elliptic Problem Solvers*, Vol. 2, Academic Press, 2014.
- [18] J.M. O'Kane, A Gentle Introduction to ROS, Independently published, 2013, Available at <http://www.cse.sc.edu/jokane/agitr/>.
- [19] J. Oyekan, H. Hu, Towards Autonomous Patrol Behaviours for UAVs, in: Proceedings of UK EPSRC Workshop on Human Adaptive Mechatronics, Stafford, UK, 2009.
- [20] Yan Yin, Sharf Inna, Forbes James, Nonlinear optimal control of holonomic indoor airship, in: AIAA Guidance, Navigation, and Control Conference, American Institute of Aeronautics and Astronautics, 2012. <http://dx.doi.org/10.2514/6.2012-4456>.
- [21] C.G. David St-Onge, N. Reeves, Dynamic modelling and control of a cubic flying blimp using external motion capture, *Proc. Inst. Mech. Eng. I* 229–210 (2015) 970–982.
- [22] O. Dugas, P. Giguère, I. Rekleitis, 6DoF camera localization for mutually observing robots, in: International Symposium on Robotics Research ISRR, Singapore, 2013.
- [23] P. Giguère, I. Rekleitis, M. Latulippe, I see you, you see me: Cooperative localization through bearing-only mutually observing robots, in: IEEE/RSJ Int. Conf. on Intel. Robots and Systems, Portugal, 2012, pp. 863–869.
- [24] M. Faessler, E. Mueggler, K. Schwabe, D. Scaramuzza, A monocular pose estimation system based on infrared LEDs, in: IEEE International Conference on Robotics and Automation, ICRA, 2014, pp. 907–913.
- [25] O. Dugas, Localisation relative à six degrés de liberté basée sur les angles et sur le filtrage probabiliste, (Master's thesis), Université Laval, Quebec City, Quebec, Canada, 2013.
- [26] A. Harmat, I. Sharf, M. Trentini, Parallel tracking and mapping with multiple cameras on an unmanned aerial vehicle, in: Proceedings of the International Conference on Intelligent Robotics and Applications, vol. 1, 2012, pp. 421–432.
- [27] A. Harmat, M. Trentini, I. Sharf, Multi-camera tracking and mapping for unmanned aerial vehicles in unstructured environments, *J. Intell. Robot. Syst.* (2014) 1–27.
- [28] M.J. Tribou, A. Harmat, D.W.L. Wang, I. Sharf, S.L. Waslander, Multi-camera parallel tracking and mapping with non-overlapping fields of view, *Int. J. Robot. Res.* 34 (12) (2015) 1480–1500.
- [29] G.F.A. Franklin, J.D.A. Powell, M.L.A. Workman, *Digital Control of Dynamic Systems*, Addison Wesley Longman, 1998.
- [30] R. Kelly, V.S. Davila, J.A.L. Perez, *Control of Robot Manipulators in Joint Space*, Springer Science & Business Media, 2006.
- [31] S. Shibata, B. Sahbi, K. Tanaka, A. Shimizu, An analysis of the process of handing over an object and its application to robot motions, *J. Guid. Control Dyn.* 1 (1997) 64–69.
- [32] R.C. Arkin, M. Fujita, T. Takagi, R. Hasegawa, An ethological and emotional basis for humanrobot interaction, *Robot. Auton. Syst.* 42 (2003) 191–201.
- [33] C. Breazeal, Emotion and sociable humanoid robots, *Int. J. Hum.-Comput. Stud.* 59 (2003) 119–155.
- [34] M. Joosse, M. Lohse, V. Evers, Lost in proxemics: spatial behavior for cross-cultural HRI, in: Proc. of the 9th Human–Robot Interface, Workshop: “Culture Aware Robotics”, 2014, pp. 184–185.
- [35] C. Bartneck, D. Kulic, E. Croft, S. Zoghbi, Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots, *Int. J. Soc. Robot* (2009) 71–81.
- [36] G. Hoffman, W. Ju, Designing robots with movement in mind, *J. Hum. Robot Interact.* 3 (1) (2014) 89. <http://dx.doi.org/10.5898/JHRI.3.1.Hoffman>. URL <http://humanrobotinteraction.org/journal/index.php/HRI/article/view/168>.
- [37] P. Brossier, Automatic annotation of musical audio for interactive systems, (Ph.D. thesis), Centre for Digital music, Queen Mary University of London, London, UK, 2006.
- [38] P. Giguère, G. Dudek, A simple tactile probe for surface identification by mobile robots, *IEEE Trans. Robotics* (2011) 534–544.
- [39] U.C. Allard, D. St-Onge, P. Giguère, F. Laviolette, B. Gosselin, Choreographic sequence detection with an electromyographic armband, *University Laval Internal Report*, 2015.
- [40] Y. Girdhar, P. Giguère, G. Dudek, Autonomous adaptive exploration using realtime online spatiotemporal topic modeling, *Int. J. Robot. Res.* 33 (4) (2013) 645–657. <http://dx.doi.org/10.1177/0278364913507325>. URL <http://ijr.sagepub.com/cgi/doi/10.1177/0278364913507325>.
- [41] Y. Girdhar, D. Whitney, G. Dudek, Curiosity based exploration for learning terrain models, in: 2014 IEEE International Conference on Robotics and Automation, ICRA, IEEE, 2014, pp. 578–584. <http://dx.doi.org/10.1109/ICRA.2014.6906913>. arXiv:1310.6767; URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6906913>.
- [42] Y. Girdhar, G. Dudek, Exploring underwater environments with curiosity, in: 2014 Canadian Conference on Computer and Robot Vision, IEEE, Montreal, 2014, pp. 104–110. <http://dx.doi.org/10.1109/CRV.2014.22>. <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6816831>.



David St-Onge Following his undergraduate studies at the École Polytechnique de Montréal, he completed a Master degree at the Robotics Laboratory of University Laval. The same project was then subject to a team grant for which he was the technical coordinator. Between these diploma and before going back to University Laval for a Ph.D. degree supervised by Prof. Clément Gosselin, he worked 7 years in the field of arts (theater, museums and circus). He is a project engineer and research professional associated with Prof. Nicolas Reeves in the NXI GESTATIO Design Laboratory since 2006. Out of interest he also completed a Master in project management and a certificate in Russian studies. Since 2016, he is a post-doctoral researcher at the MIST Laboratory with Prof. Giovanni Beltrame.



Pierre-Yves Brèches Currently an IT consultant at Bosch Engineering, Abstatt, Germany. He completed a master with a double diploma cursus between École Centrale de Lyon and McGill University in Mechanical engineering. He completed his undergraduate studies in France at École Centrale de Lyon and continued onto graduate studies in the Aerospace Mechatronic Lab of McGill.



Inna Sharf Currently a full professor in the Department of Mechanical Engineering at McGill University, Montreal, Canada. She received her B.A.Sc. in Engineering Science from the University of Toronto (1986) and her Ph.D. at the Institute for Aerospace Studies, University of Toronto (1991). Prior to relocating to McGill in 2001, she was on faculty with the Department of Mechanical Engineering at the University of Victoria. Sharf's research activities are in the areas of dynamics and control with applications to space robotic systems, unmanned aerial vehicles and legged robots. Sharf has published over 100 conference

and journal papers on her academic research. She is an associate fellow of AIAA and a member of IEEE and ASME.



Nicolas Reeves Architect, Artist, Researcher. He studied in France, Canada and USA. He was Vice-President of the Society for Arts and Technology (03–08), Scientific Director and then director of the Hexagram Institute for Research/Creation in Media Arts and Technologies (05–15). He is a Full Professor at the School of Design, University of Quebec in Montreal. He also directs the NXI GESTATIO Design lab (computer, arts, architecture). His education is in Physics (U. of Montreal, 1986) and in Architecture (MIT, 1988). He trained at Atelier Philippe Madec (Paris) and the Jacques Rousseau studio (Montreal). His artistic research explore the basis of the concepts of order, organization, and information.



Ioannis Rekleitis Currently an Assistant Professor at the Computer Science and Engineering Department, University of South Carolina, and an Adjunct Professor at the School of Computer Science, McGill University. In 2004–2007 he was a visiting fellow at the Canadian Space Agency working on Planetary exploration and On-Orbit-Servicing of Satellites. During 2004 he was at McGill University as a Research Associate in the Centre for Intelligent Machines with Professor Gregory Dudek in the Mobile Robotics Lab (MRL). Between 2002 and 2003, he was a Postdoctoral Fellow at the Carnegie Mellon University in

the Sensor Based Planning Lab with Professor Howie Choset. He was granted his Ph.D. from the School of Computer Science, McGill University, Montreal, Quebec, Canada in 2003. His Research has focused on mobile robotics and in particular in the area of cooperating intelligent agents with application to multi-robot cooperative localization, mapping, exploration and coverage. He has done extensive work on space, aerial, ground, surface, and underwater robots, with more than 60 journal and conference articles. His interests extend to computer vision and sensor networks.



Patrick Abouzakhm He completed a master at the Aerospace Mechatronics Lab at McGill University on the docking of Tryphon blimps. He received his bachelors degree in Mechanical Engineering from McGill and has continued onto graduate studies with a concentration in Dynamics and Control under the supervision of Prof. Inna Sharf.



Yogesh Girdhar Now a Devonshire Postdoctoral Scholar, and a FQRTN Postdoctoral Fellow at the Woods Hole Oceanographic Institution, in the Applied Ocean Physics and Engineering department. He received his BS and MS from Rensselaer Polytechnic Institute in 2005, where he was awarded Paul A. McGloin prize for most outstanding academic achievement in Computer Science. He received his Ph.D. from McGill University in 2014 with thesis title Unsupervised Semantic Perception, Summarization, and Autonomous Exploration for Robots in Unstructured Environments, which received the Honorable Mention for CIPPR Doctoral Dissertation Award. Yogesh's research interest is in algorithmic problems that lie at the intersection of field robotics, machine learning and computer vision.



Adam Harmat He completed his Ph.D. at McGill University working on autonomous takeoff, flight, and landing behaviors for a small multirotor aircraft using mostly vision-based methods. He has a formal background in mechatronics engineering, but he is also extremely interested in problems in computer vision and machine learning.



Gregory Dudek Professor of computer science at McGill University, he was the Director of the McGill Center for Intelligent Machines from 2004 to 2007, and he is the current Director of the McGill University School of Computer Science. He was made a Dawson Scholar of that university and directs the Mobile Robotics Laboratory there. He was awarded a James McGill chair in 2008. He has written over 150 refereed articles on computer vision and robotics, and is co-author (with Michael Jenkin) of the major book Computational Principles of Mobile Robotics. His lab is among the first in Canada to deal substantively and specifically with issues concerning mobile robotic systems.



Philippe Giguère Professor at the Department of Computer Science and Software Engineering at Laval University (Quebec City) since 2010. He has 12 years of experience in mobile robotics and sensing in academia, and 6 years of experience in embedded, real-time systems in the private industry. He leads a research group in mobile robotics and artificial perception within the DAMAS research laboratory. The main focus of his research is to increase the autonomy of intelligent robotics systems deployed in challenging unstructured environment (underwater or natural), through the application of advanced machine learning techniques or sensor fusion. His most recent projects were related to terrain identification, tactile perception, computer vision (2D/3D), localization, signal processing (speech) and the use of terrestrial LiDAR. He is also a member of the FRQNT-REPARTI network on intelligent environments, the center for big data of Laval University and distal member of the NSERC Canadian Field Robotics Network.