

Humanoid Robots in Aircraft Manufacturing

The Airbus Use Cases

By Abderrahmane Kheddar, Stéphane Caron, Pierre Gergondet, Andrew Comport, Arnaud Tanguy, Christian Ott, Bernd Henze, George Mesesan, Johannes Engelsberger, Máximo A. Roa, Pierre-Brice Wieber, François Chaumette, Fabien Spindler, Giuseppe Oriolo, Leonardo Lanari, Adrien Escande, Kevin Chappellet, Fumio Kanehiro, and Patrice Rabaté



©ISTOCKPHOTO.COM/FATIDIO

We report on the results of a collaborative project that investigated the deployment of humanoid robotic solutions in aircraft manufacturing for several assembly operations where access by wheeled or railported robotic platforms is not possible. Recent developments in multicontact planning and control, bipedal walking, embedded simultaneous localization and mapping (SLAM), whole-body multisensory task-space optimization control, and contact detection and safety suggest that humanoids could be a plausible solution for automation, given the specific requirements in large-scale manufacturing sites. The main challenge is the integration of these scientific and technological advances into two existing humanoid platforms: the position-controlled Human Robotics Project (HRP-4) and the torque-controlled robot (TORO). This integration effort was demonstrated during a bracket-assembly operation inside a 1:1-scale A350 mockup of the front part of the fuselage at the Airbus Saint-Nazaire site. We present and discuss the main results achieved in this project and provide recommendations for future work.

Digital Object Identifier 10.1109/MRA.2019.2943395

Date of current version: 30 October 2019

Collaborative Humanoid Robots

On 21 February 2019, for the first time, two humanoid robots, the National Center for Scientific Research (CNRS)-Laboratory for Computer Science, Robotics, and Microelectronics of Montpellier's position-controlled HRP-4 and German Space Agency's TORO, accessed the Airbus civilian airliner manufacturing plant at Saint-Nazaire, France, and achieved the final demonstration of the European Union's collaborative project Comanoid [30] (i.e., collaborative humanoid robot), which gathered four academic European partners, a Japanese national research institute, and the end-user Airbus. The goal of Comanoid is the eventual feasible and plausible deployment of humanoid robotic technology as an automation solution to achieve specific non-added-value tasks in aircraft manufacturing operations.

Comanoid focuses on obtaining precise accessibility (i.e., for areas where wheeled robots cannot be deployed) through whole-body, multicontact planning motion with advanced, embedded, 3D dense SLAM localization and visuoforce servoing capabilities under safety constraints. Another CNRS-National Institute of Advanced Industrial Science and Technology, Joint Robotics Laboratory and Airbus bilateral project (not discussed in this article) complements Comanoid with more complex manipulation tasks, such as fast

torqueing, flexible cable assembly, cleaning, and cockpit assessment operations.

Despite the considerable advances in many fundamental and technical aspects of humanoid robotics research, public and private subsidies have become more demanding of the motivations and perspectives behind humanoid research, even explicitly seeking so-called “killer” application(s). Aside from pure robotics, humanoid robotics can be useful in research related to embodiment and consciousness and, more generally, in neuro and cognitive sciences. They are also envisioned as receptionists in entertainment sectors and in various social interactions, e.g., as “assistants” for frail persons (note that not all of these applications require having robots with legs); this is part of the business plan for companies such as SoftBank Robotics in Japan and PAL Robotics in Spain. Other services, such as disaster emergencies, nuclear power plant dismantling, space exploration, and space extravehicular activities, also consider humanoids, to some extent, as remote-intervention robots. Yet no industrial sector has expressed a need in terms of humanoid robots for manufacturing. The exception is the Glory factory in Saitama, Japan, where the assembly of money-handling machines and the like incorporates humanoid (torso) robots (Nextage) for manufacturing.

Our work is the result of many years of discussions and exchanges with Airbus Group Innovations about what solution is suitable for robotic automation of non-added-value tasks. During preliminary discussions, humanoid robots were not considered because of their complexity, nonavailability, and slow progress. The next section briefly discusses the needs in automation that might call for a bipedal solution.

Why Humanoid Robots in Aircraft Manufacturing Automation?

On 26 February 2014, a workshop was organized on the premises of the French embassy in Tokyo. The workshop gathered academic experts in robotics to complement on-going Airbus internal investigations concerning the best robotic technology for automating non-added-value tasks (for Airbus workers), such as measuring and positioning (printing) the bracket templates in the cargo hold and other hard-to-access areas. Initial inquiries suggested that a legged robot carrying manipulator arms could be one solution. Then, the other question was: How many legs/arms?

At first, humanoid technology was not considered because it was thought that two-legged robots were not reliably stable. Therefore, hexapods or quadruped Centaur-type robots were considered. Nonetheless, multi-legged solutions come with some practical limitations and drawbacks.

1) The stability benefits of a multilegged robot’s (and even a wheeled robot’s) larger support area are reduced or lost when the mounted manipulators approach or go beyond the support borders.

2) Indirectly related to 1), multilegged robots occupy a larger shop-floor space, which could not be feasible in many areas of the aircraft because of confined, cluttered, or narrow spaces (e.g., human workers crawl to access the cockpit area).

3) The cost is higher for multilegged robots porting dual-arm manipulators.

In light of these limitations, and with demonstrations of HRP-2 humanoid multicontact technology that showed a humanoid can also use its other limbs for contact support (as humans do), we decided with Airbus to investigate further humanoid technology solutions.

Context and Automation Needs

The aeronautical industrial context in terms of the final assembly line and major component assembly can be best explained through the following specificities:

- *Large structures and millions of parts:* For example, the A380 has more than 1 million rivets, and multiple dozens of thousands of supports must be assembled (such as for cables, tubes, brackets, and so on).
- *Low production rate:* For the A320, it is 1.5 airplanes per day; for the A380, it is four per month. This rate could increase if production delays are avoided.
- *Very high quality requirements:* For example, hole positioning requires sub-millimeter precision, and a hole-normality adjustment requires less than 1° of clearance.
- *Skilled operators:* Complex and cognitive high-value tasks are allocated for skilled operators, freeing them from simple, repetitive, and heavy non-added-value tasks.
- *Workspace:* Centered around the inside and fuselage with limited space, the workspace is generally a crowded area.

The arrival of robots in aeronautic manufacturing plants is more recent (the 1990s) due to the low accuracy of industrial robots mostly designed for the automotive industry. Initial applications consisted of drilling holes on medium-sized components, such as nose fuselage structures. Until recently, there were no automated means on the final assembly line, except for the circumferential junction of A340 automation; however, this was abandoned because it was not flexible enough.

For these reasons, Airbus sought collaborative robots capable of performing assembly tasks inside aircraft. A few years ago, the Airbus Puerto Real plant conducted the first experiments using Nextage Kawada’s robot to perform simple assembly tasks. To our knowledge, apart from this first experiment, no real work has been achieved

At first, humanoid technology was not considered because it was thought that two-legged robots were not reliably stable.

and published on the use of humanoid robots in the aeronautical industry.

The main concerns are non-added-value tasks because they include health risks, such as high-repetitive strain injury, that require high precision and low dexterity. With Comanoid, we targeted the tasks of pattern printing or bracket positioning on the fuselage because they are made nearly everywhere and do not require heavy payload nor dexterous manipulation. The primary goal was to show the humanoid robot's ability to reach the workspaces and for it to be precise with both the localization and positioning of the end effector in the structure.

In the framework of another research program, four use cases of relative complexity were identified:

- 1) circuit breaker automation, which requires operation in cluttered spaces
- 2) hydrofuge protection and cleaning, which requires managing equilibrium and in-line tool trajectory planning
- 3) torquing automation, which requires multicontact balance control in dynamic motion and force control
- 4) system installation, which requires manipulating flexible objects.

Other tasks, such as drilling and cockpit end-phase systems checks, were also suggested.

Constraints

Accessibility and manipulation needs are already very challenging. The additional constraints under which robots operate render these challenges nearly unreachable, as depicted in Figure 1.

- Any robot shall operate and share the same space with human workers and be safe in all circumstances.
- Humanoid falling shall be demonstrated in worst-case conditions on humans and even on airplane structures.
- The robot shall be able to access any spot starting from any other spot and travel between the three working levels using human means, i.e., without requiring changes to the manufacturing infrastructure.
- No external sensors, such as external cameras, are allowed; however, bringing additional lighting is permitted.
- Safety certification and compliance with existing internal regulations and procedures will be established.

The good news is that workers follow strict procedures and rules. The CAD model of any piece or part of the airplane

is available (including its exact shape and inertia parameters), and a digital mockup is updated in parallel to the entire assembly process.

Basic Technological Requirements

In light of previously established requirements and needs, we have identified the following must-have technologies to be developed further for the Comanoid project and implemented in the HRP-4 and TORO humanoid robots by the end of the project: 1) multicontact planning and control, 2) walking, 3) embedded localization and mapping, 4) integrating visual and force control into the task-space optimization control, and 5) contact detection and safety.

Multicontact Planning and Control

Workers often operate in very unergonomic postures, as shown in Figure 1. In such cases, they use multicontact postures to relax stress, improve their equilibrium, or cast their body posture (e.g., contact with the knees, shoulders, and back). Humanoids shall be endowed with similar multicontact behaviors, which is a key technology of Comanoid; this allows for the transformation of a humanoid robot into a reconfigurable multilimb system that can adapt to narrow spaces and increase its equilibrium robustness and manipulation capabilities.

Multicontact Planning

A recent review of multicontact technology [1] reveals that this problem is consensually approached through the use of a multilevel computation. The first level plans contacts around a sort of free-motion preplanned "guide" that exploits some properties under different hypotheses (there are many variants) but has as an output a set of contact sequences and associated transitions. In a second phase, the latter are inputs for a simplified model [e.g., center of mass (CoM)] that generates a consistent centroidal dynamics trajectory under balanced criteria (presented later in this article). In the last phase, this generated trajectory is the input to the whole-body controller, which also deals with other task objectives and constraints. The problem in such a multilevel computation is to make each phase potentially feasible for the upcoming one and ensure that, if any phase turns out to be infeasible for the given input, the latter is quickly recomputed from the current state. Although

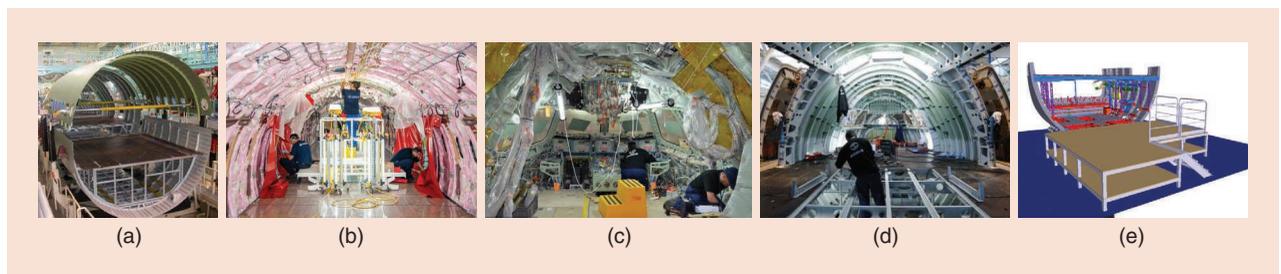


Figure 1. Examples of the working areas and conditions in aircraft manufacturing sites. (a) The cargo area. (b) Kneeling and using a small ladder to access the roof. (c) The cockpit area. (d) The removal of the shop floor to access in-between levels. (e) A CAD digital model at a 1:1 scale of the A350 mockup used to demonstrate the technologies for most of these scenarios.

impressive results were obtained, this approach does not work well in practice. In the context of manufacturing, it is certainly not a good idea to plan contacts as if we do not know where and how they should be made. Assembly operations are quite repetitive in many aspects, and only a few variations are to be dealt with locally.

Our approach is to exploit knowledge about how humans operate in similar circumstances and to consider that pre-defined procedures dictate grossly what contact sequence is necessary to conform to both safety and task needs. Then, the posture the robot takes to achieve a given task in a multicontact setting is computed, including task forces [2]. The force required for a given task has great influence on the configuration to be taken. Moreover, the task to be completed would suggest additional contacts to be established, not only to increase the equilibrium robustness but also to increase the operational force by creating additional kinematic loops [3] and exploiting internal force distribution.

Upon finding the contact configuration suitable for each operation, special attention is required for the realization of transitions between different multicontact poses. To achieve dynamic contact transitions and avoid disadvantageous quasi-static motions, we adapted the divergent component of motion (DCM) framework from dynamic walking to the generation of multicontact locomotion tasks [4]. The generated DCM and CoM trajectories are then compatible with a given sequence of multicontact poses. In addition to the CoM motion, the timing of the end-effector motions, i.e., the stance durations and contact transition timings, are optimized.

Multicontact Balance and Control

The multicontact planning phase generates desired contacts and a CoM to track at best (or exactly) by the whole-body controller, among other task objectives (such as manipulation). HRP-4 is controlled using task-space objectives in the sensory space, formulated as a weighted and constrained quadratic program. It includes visual and force servoing (see the “Visual and Force Control” section) and various task templates to achieve, sustain, or remove a contact and achieve force admittance control [5] and so on.

The general balance condition is written in a bilinear form in terms of CoM positioning and its acceleration and contact forces. For bounded convex CoM positions, we can obtain feasible CoM accelerations as a 3D cone. For bounded convex CoM accelerations, we can acquire feasible CoM positions as a 3D convex hull [6]. Resultant geometrical shapes can be used and integrated in both planning and control [1].

Multicontact tasks with TORO are achieved with a passivity-based control framework [7], [8]. The following four steps are summarized within a single, force-based optimization problem:

- 1) the realization of a desired wrench that stabilizes the CoM at a desired equilibrium
- 2) the implementation of desired impedance forces for the end effectors that perform manipulation tasks

- 3) distribution of the resulting overall wrench to the wrenches acting at the available end effectors in contact via a “grasp” matrix

- 4) the realization of all end-effector wrenches via the robot’s joint torques (see [7]–[9], where the passivity properties are analyzed).

Passivity is of utmost importance whenever the robot must interact with an unknown but passive environment. This control framework has been successfully applied to several multicontact configurations related to the given aircraft manipulation use cases, as illustrated in Figure 2.

Walking

No current humanoid robot is able to walk around an Airbus manufacturing facility with the same reliability as a human worker. Because this environment is shared with human coworkers, reliability is also a major aspect of robot and coworker safety as well as the airplane’s structure. The reliability and safety of humanoid robots are complex issues, for which no meaningful quantitative measure currently exists. As a result, and because the field of legged locomotion is still maturing, the Comanoid consortium’s strategy has been to develop reliability and safety for a diverse range of approaches to walking-motion generation and control, instead of committing to a single one.

The modeling of legged locomotion relies on a few simple physical principles [10]. The Newton equation of motion makes it clear that the robot needs contact forces f_i to move its CoM c in a direction other than that of gravity g , while the Euler equation of motion makes it clear that the position of the CoM with respect to (w.r.t.) contact points s_i is critical for keeping the angular momentum L of the robot around its CoM under control:

$$\begin{cases} m(\ddot{c} - g) = \sum f_i \\ \dot{L} = \sum (s_i - c) \times f_i \end{cases} \quad (1)$$

where m is the mass of the robot. These are the elements of motion used to maintain balance, but how? Ensuring that the robot is always able to reach a cyclic

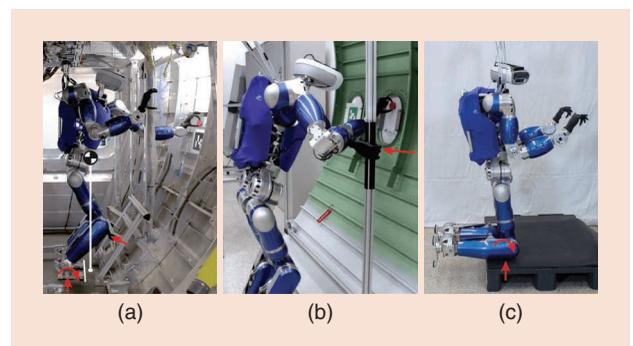


Figure 2. The different multicontact situations. (a) Feet plus knee support, (b) hand support, and (c) full kneeling.

motion or an equilibrium point in a few steps is a simple and effective way to guarantee that it is able to avoid falling. This also provides a simple method to detect the risk of an imminent fall to trigger fall-mitigation behaviors when appropriate.

The Comanoid consortium adopted a standard multi-stage framework for walking motion generation and control: 1) a short sequence of step positions and phase durations is proposed depending on robot current and goal states as well as the environment (i.e., obstacles and people); 2) the resulting motion of the CoM and contact forces are computed, making certain that a cyclic motion or an equilibrium point can be reached within a few steps; and 3) the robot is controlled to realize the computed CoM motion and contact forces.

We devised several approaches within this common framework to improve reliability and safety. Some address stages 1 and 2 in the previous paragraph separately [11], while other approaches do so in a single stage [12]. This is slightly more involved numerically, but it improves the robot's capacity to react effectively to a dynamic environment. In static environments, stages 1 and 2 can be considered once every few steps. Otherwise, e.g., with workers moving around, these stages are reevaluated more often, using approaches based on model predictive control (MPC).

Capturability approaches guarantee that the robot is consistently in a viable state and is always able to stop safely [12]. Alternative approaches investigate the more general concept of boundedness [13]. One approach that integrates safety guidelines w.r.t. surrounding humans is to adapt the current goal of the robot w.r.t. the current state of the environment [14]. One can also integrate collision mitigation and passive safety constraints directly in a combination of stages 1 and 2 [12]. This is more involved numerically, but it improves the robot's capacity to navigate safely in the presence of workers.

To negotiate uneven ground and stairs, the robot's CoM height should be adjustable, thus introducing the nonlinearities present in stage 2. These nonlinearities are sometimes neglected, but at the risk of failure. We handle them explicitly by considering a piecewise, linear 3D trajectory of the DCM [11], or we bound them by constraining the height variations of the CoM above the ground, adapting its capturability and boundedness accordingly [15].

This diversity of approaches has been used to demonstrate two different humanoid robots capable of navigating in a typical Airbus environment, walking reliably to their exact destinations to complete effectively the assigned manipulation tasks. This requires a tight integration of walking with visual servoing and SLAM, multicontact phases, and safety guidelines. This tight integration is arguably the biggest achievement of the project, as walking is not addressed independently from navigation, manipulation, and safety issues.

In-Site and In-Craft Localization and Mapping

Autonomous SLAM in aircraft manufacturing is a fundamental capability for practical use of humanoid robots in a real-world setting. Few real-time (RT) approaches have been proposed that can account for dynamic environments and long-term incremental changes (e.g., the manufacturing process and lighting variation) while maintaining positioning accuracy or tolerating loss of tracking. Moreover, assembly operations require continuous correspondence between the evolving digital mockup and the reality of the airplane assembly. Consequently, semantic knowledge is highly necessary.

Image-based keyframe navigation was used for its efficiency and accuracy in humanoid positioning because it allows for closing the feedback control loop in the sensor space, subsequently avoiding drift, improving robustness, and enabling loop closure and relocalization (see the results in the "Integration and Experiments" section).

Keyframe SLAM and Image-Based Navigation

A direct, multi-keyframe approach is used to perform red-green-blue-depth (RGB-D) SLAM for navigation [16]. The sensor pose $\xi \in se(3)$ is estimated w.r.t. the set of closest keyframes by minimizing the error between the current frame and a predicted one. The current measurement vector for each intensity and depth i is defined as $\mathbf{M}_i = [\mathbf{P}_i^T \mathbf{I}_i]^T \in \mathbb{R}^4$. The predicted keyframe $\mathbf{M}_i^* \in \mathbb{R}^4$ is obtained by blending the n closest keyframes (typically five) at the last pose estimate. The point-to-hyperplane iterative closest-point approach [17] is then used to estimate the pose iteratively from the following error function:

$$\mathbf{e}_i(\xi) = \mathbf{N}_i^{*T}(\mathbf{M}_i^* - w(\mathbf{M}_i, \xi)), \quad (2)$$

where the normals $\mathbf{N}_i^{*T} \in \mathbb{R}^4$ are computed once on the referenced 4D measurement vector and $w(\cdot)$ is the warping function that transforms the current image to the reference, based on the current pose estimate.

This basic alignment procedure is extended to large environments by using a keyframe graph, built and refined incrementally as mapping is performed. To take advantage of the topometric keyframe representation, the target position is given as a sensor-based keyframe to reach. This permits the direct image-based error defined in (2) to be minimized by the robot controller, effectively allowing the robot to position itself with high accuracy locally while tolerating global drift.

Map Reuse, Place Recognition, and Time-Varying Environments

Due to the complexity of the manufacturing site's factors, such as workers, tools, and assembly increments, a prior keyframe graph was created from a static environment and then reused online. To ensure robot positioning, pose uncertainty and sensor-based errors were monitored to determine tracking loss and trigger relocalization (i.e., pose estimation w.r.t. the nearest keyframes). In addition, the robot's viewing direction

was actively controlled to observe the closest keyframes that improve robustness.

In practice, partial keyframe maps were acquired for each task workspace of the robot. These submaps were connected only topologically. They were then reloaded online according to the robot's location. During the demonstrations, this proved to be very efficient because there were even more people in the environment than in a typical work setting.

Semantic RT Mapping

Planning and control algorithms require higher-level knowledge about the environment surrounding the robot. A rich source of information for learning this knowledge is the Airbus digital mockup. An RT semantic segmentation network was developed where classes are learned from the labeled digital mockup. The segmentation network is fine-tuned for the given use case and trained in a semisupervised manner using noisy labels. The method is optimized for RT performance and integrated with the Robotic Operating System (ROS) to provide semantic reconstruction for navigation.

Semantic information can also be fed back to the SLAM. The measurement vector \mathbf{M} is extended to simultaneously minimize photometric, geometric, and semantic costs in RT [18]. The approach has been shown to robustly construct a labeled, large-scale 3D map relevant to robotic tasks. In situ experiments at Airbus validated our approach, as depicted in Figure 3.

Visual and Force Control

Visual Tracking and Servoing

For accurate manipulation, the amount of observed data is not sufficient to provide its localization through the SLAM described previously. For this reason, particular object detection, visual tracking, and visual-servoing techniques have been developed in Comanoid. The object detection is done using classical keypoint matching methods and provides an initial localization (full pose) between the observed object and the vision sensor. The object and its successive poses are then tracked during the humanoid motion thanks to a modular, RT, model-based vision tracker (MBT). This MBT is available for the community in the latest release of ViSP [19]. It can combine low-level edges extracted in the image, textured point of interest, and depth [20] when an

RGB-D camera is available, which is the case for both TORO and HRP-4.

Once visual data are available at a high rate (10 Hz, in our case), they can be embedded in a closed-loop visual-servoing scheme. However, to be combined with the quadratic programming (QP) controller of the HRP-4, it has been necessary to go from the classical kinematics modeling of the visual features to their dynamic modeling. This modeling step is described in [21]. It is then possible to consider tasks such as gaze control or set-point reaching, occlusion avoidance, or preserving the visibility of an object, as any other robotics task and constraint.

The object detection, visual tracking, and visual servoing described previously have all been implemented in TORO (see Figure 4) and in HRP-4 (see Figure 5). Note that with

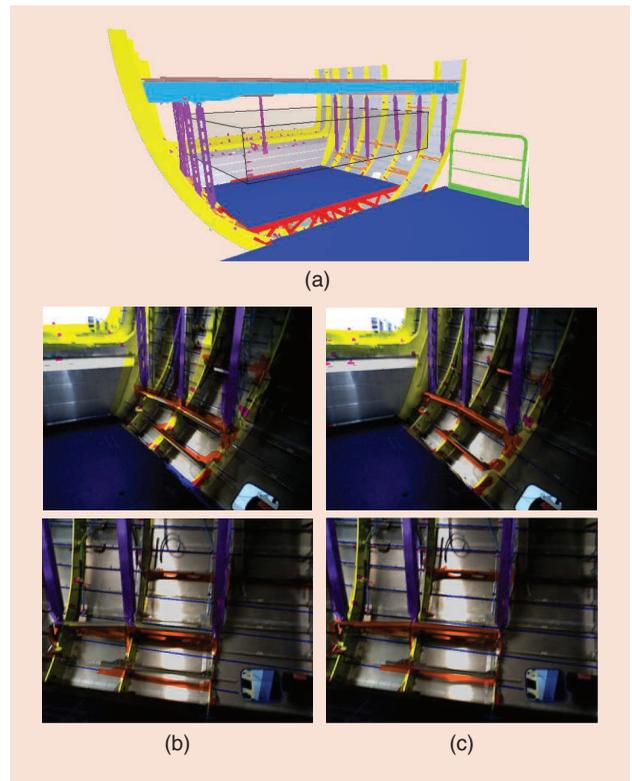


Figure 3. Overlays of RGB images and their ground-truth labels projected from (a) a synthetic 3D model of the aircraft mockup, (b) initially using global registration, and (c) after postrefinement using an iterative closest-point algorithm.

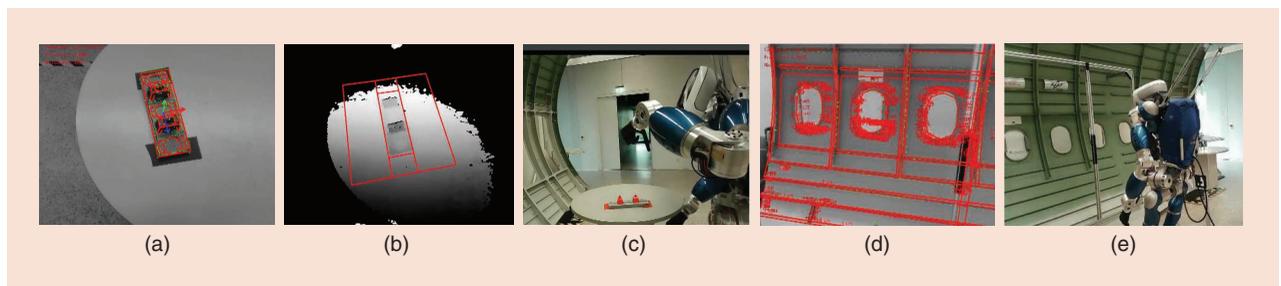


Figure 4. The visual tracking results obtained using TORO. (a) An image with the edges extracted and a model of the bracket projected from the estimated pose. (b) A depth map with the bracket model projected from the same pose. (c) A view of TORO and the tracked bracket put on a table. (d) Fuselage tracking using edges and keypoints. (e) TORO moving toward the fuselage using visual tracking.

HRP-4, tracking both the object of interest and the HRP-4 left-arm gripper is done simultaneously, which provides a robust and accurate positioning of the gripper w.r.t. the object to be manipulated.

Force Control

Force control is necessary in many operations, such as pulling the circuit breakers and gluing the bracket on the fuselage. To complement vision, force control is implemented on both robots. Whereas force control on TORO is rather straightforward due to its torque control capabilities, it requires specific considerations for the position-controlled HRP-4.

Force control is necessary in many operations, such as pulling the circuit breakers and gluing the bracket on the fuselage.

First, with contact or interaction forces being the decision variable of the QP, constraints such as nonsliding linearized friction cones, bounds on applied forces (in any direction) that might be dictated by the task, or the minimal (threshold)

force required to sustain a contact can easily be added. However, to servo a given terminal point to a force, we must decide upon that desired force. If the desired force is user given, it can be realized by an admittance task, which is part of the QP formulation [5]. Yet defining arbitrary forces to hold a contact may unnecessarily restrict the motion of the robot, especially in a multicontact setting. For instance, we want to apply with the left gripper a force of roughly 10 N in the normal direction of the fuselage surface to glue the bracket, but we would rather avoid specifying a particular force for the right gripper leaning on the structure to support the motion. In this case, we simply require that force be greater than a given threshold [5].

Safety and Contact Detection

Safety is critical in human-robot collocated spaces, such as in aircraft manufacturing, but it is also important to

preserve the integrity of the humanoid itself as well as the surrounding environment. Although the problem of safe coexistence with humans had been widely addressed for fixed-based manipulators, there were no results for humanoids before Comanoid.

Designing robust walking controllers that perform reliably is already one approach to increase the intrinsic level of safety. As a more systematic, complementary approach, we have identified a set of safety guidelines [22] from which it is possible to derive several safety behaviors, grouped in three categories. Override behaviors (e.g., emergency stop) will stop the execution of the current task and lead to a state from which normal operation can only be resumed by means of human intervention. Temporary override behaviors (e.g., evasion [14]) will also suspend task execution but only for the limited amount of time needed to handle the safety concerns, after which task execution is automatically resumed. Finally, proactive behaviors (e.g., human visual tracking or footstep adaptation) do not stop the task but rather attempt to increase the overall safety level by calling for a modification of the current robot activity. The activation of the safety behaviors, integrated in an MPC control framework, is orchestrated by a suitable state machine based on the robot's on-board sensor feedback.

Safe falling strategies are important for minimizing robot damage in case of loss of balance. Our approach for the HRP-4 humanoid aims to actively reshape the robot toward one favorable configuration from a set that has been identified in advance [23]. Control gains are then adapted in RT to comply with its postimpact dynamics. For the TORO robot, the use of passive protections (e.g., pads and airbags) has been found to significantly reduce (by more than 50%) the accelerations experienced by robot linkages during the fall. Moreover, an active falling strategy based on energy minimization can be adopted to minimize impact velocities at the elbows. Note that active falling strategies can be considered and realized as override behaviors in the aforementioned safety framework.

Intentional or accidental contact detection is currently embedded in some commercial fixed-base manipulators, while contact localization and force reconstruction were

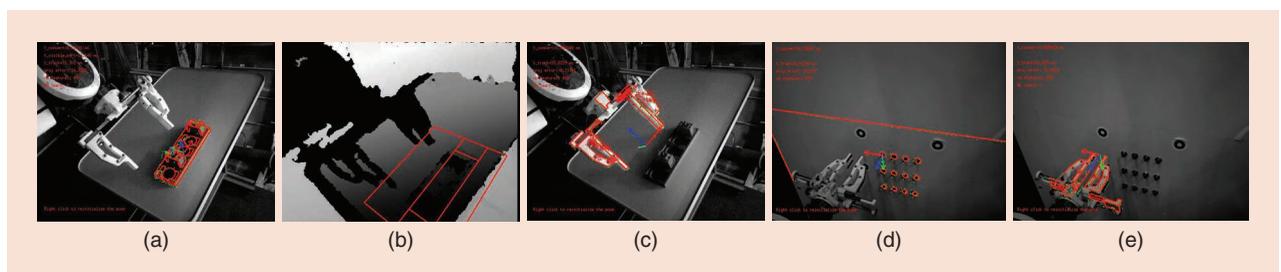


Figure 5. The visual tracking results obtained with HRP-4. (a) An image with the edges extracted and a model of the bracket projected from the estimated pose. (b) A depth map with the bracket model projected from the same pose. (c) An image with the edges and keypoints used to track the gripper, and a gripper model projected from the estimated pose: a circuit breaker panel tracking. (d) An image with edges and a model of circuit breakers projected from the estimated pose. (e) The same image with edges and a gripper model projected from the estimated pose.

achieved in prototypes using external depth cameras; again, no results for humanoid robots were available before Comanoid. A number of methods that solve this problem were developed during the project, either based on the mismatch of expected versus measured torques [24] or by using a nonlinear observation of generalized momenta [25]. Successful experimental demonstrations were achieved on HRP-4 humanoid.

In highly cluttered environments, contacts might occur frequently. Sometimes, when unexpected or undesired contacts are detected, it is already too late to react, and replanning might be required. Therefore, the idea of having a whole-body, low-range distance field sensing could be an appealing functionality. We have investigated capacitive sensing because it is widely used and proved to be efficient in robotics, see, e.g., Fogale [31], with which we paired to challenge 1) customizing the technology to whole-body humanoids and 2) increasing the sampling frequency to include the signals in the control-loop. Because we did not have spare HRP-4 covers, we had to build inner covers to hold the electrodes for each link, as presented in Figure 6. There is a total of 17 subshells in the HRP-4: one each in the head, torso, back, front and back waist; two per shoulder; two per upper arm; and two per upper leg, each equipped with a total of 54 self-capacitive electrodes that can statically measure the electric capacitance formed between a surrounding conductive object coupled to the ground and itself. It transmits collected measurements in one User Datagram Protocol frame at 1 kHz. We experienced difficulties during final assembly because the added wires that needed to be routed inside the HRP-4 structure,

following the existing ones and the additional thickness (although of millimeter order), made the mounting back of the covers tighter.

Integration and Experiments

We report on the final demonstrator of the Comanoid project defined by the end user as follows: 1) the robot is at floor 0, 2) it positions itself in front of stairs, 3) it climbs the stairs to the first floor, 4) it reaches the aircraft demonstration area, 5) it grasps the parts or tools disposed of on top of a table, 6) it moves into the predefined working area, 7) it performs predefined tasks accurately, 8) it goes out of the aircraft, and 9) it returns to floor 0.

Because of its hardware limitations and logistics, as well as to meet the limitations of time and space, each humanoid robot performed the maximum tasks possible with some variants. The end user suggested using the bracket's positioning on the fuselage.

The demonstrations could have been performed on the real aircraft, but it was not a requirement for the project, as entering the production line requires specific approvals that are difficult to obtain at this developmental stage. A physical mockup of the A350 section at a 1:1 scale, representative of the actual airplane section, is available at the Airbus site. It is used to

The mc_rtc framework is written in C++ but permits the writing of any robot controller in either C++ or Python.

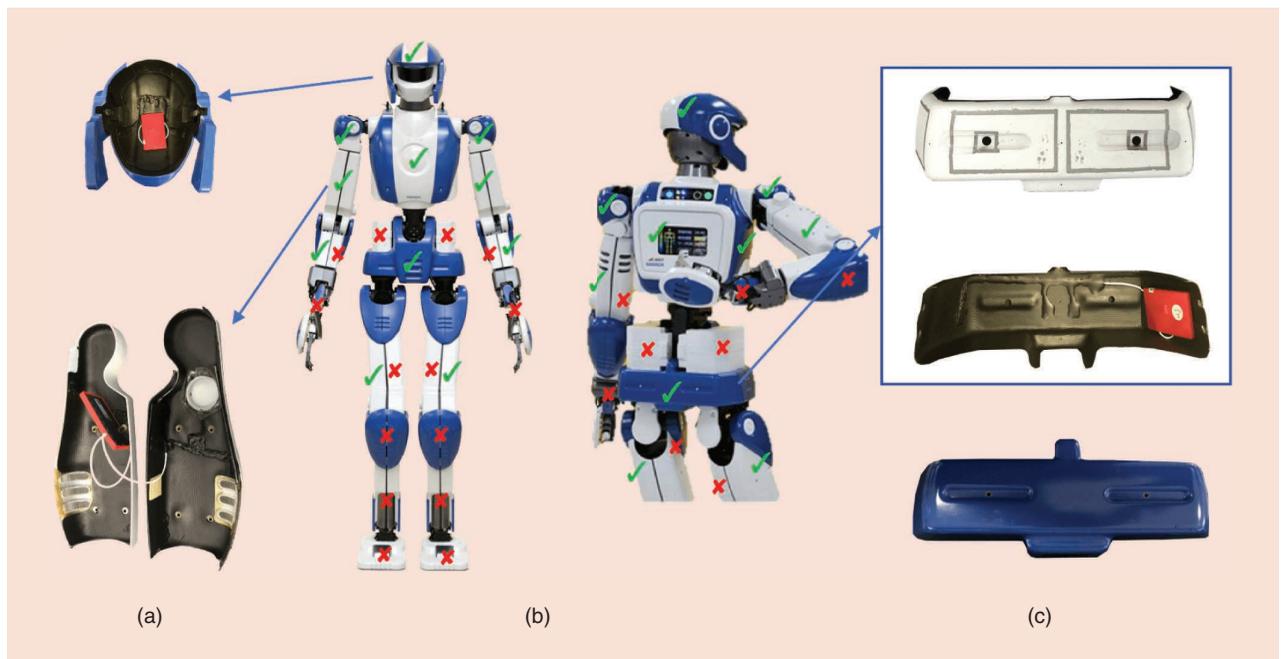


Figure 6. Fogale's capacitive sensing customized for HRP-4. (a) The head and one arm fully equipped/mounted with the capacitive-sensing solution in the embedded electronics. (b) The sensitive areas. (c) A subshell split; the upper part is the back/face of the intermediary cover hosting the electrodes marked with a gray rectangle, which is mounted on the original cover (in blue).

assess new technological ideas in production and assembly operations; however, a significant number of parts were machined from metal blocks instead of being assembled from production parts. Given that the A350 has a significant number of composite parts, the materials and textures might differ from the aircraft even if the partial geometry is very close.

Demonstrator of the HRP-4 Humanoid Robot

The demonstrator of the HRP-4 humanoid robot has the following variants w.r.t. the previously described scenario:

- Localization everywhere uses SLAM solely.
- Climbing/walking uses its own developed controller.
- Stair climbing is accomplished without the use of handrails.
- Task-aware, multicontact planning uses the left hand for additional support to lean toward the structure.
- Localization (also grasping) of the bracket and the task-aware contact uses visual servoing.
- The bracket is glued and released with force control.
- Safety achieved using whole-body capacitive sensing is demonstrated interactively after the robot exits the working area.
- All of the demonstrator's tasks are completed autonomously by the robot, but the operator validated each step prior to transitioning to the next (this can be skipped) and intervened if SLAM relocalization failed.

- Two full trials are performed without securing the robot.
- HRP-4 did not return to floor 0.

Figure 7 illustrates the main components of the task specification and control architecture for control of the HRP-4 humanoid robot, using the technological bricks described in the “Basic Technological Requirements” section. This architecture is currently implemented in several other robots, such as the Soft-Bank Robots Pepper and Nao, other HRP robots, and the Sawyer robot. It has three main components:

- 1) low-level and high-performance C++ libraries for robotic experts
- 2) a unified controller interface `mc_rtc`, which is the control framework used to facilitate the development of controllers and the integration of new robots
- 3) simulation/control interfaces that are “glue layers” between `mc_rtc` and simulation, e.g., virtual robot experimentation platform (V-REP), Choreonoid, or robot-hardware interface.

The `mc_rtc` framework is written in C++ but permits the writing of any robot controller in either C++ or Python. Our low-level libraries are mainly concerned with the mathematical and numerical aspects of control, i.e., computing all of the required matrices and vectors correctly in a timely manner and setting up and solving optimization problems. The `mc_rtc` framework brings simpler interfaces, simpler semantics

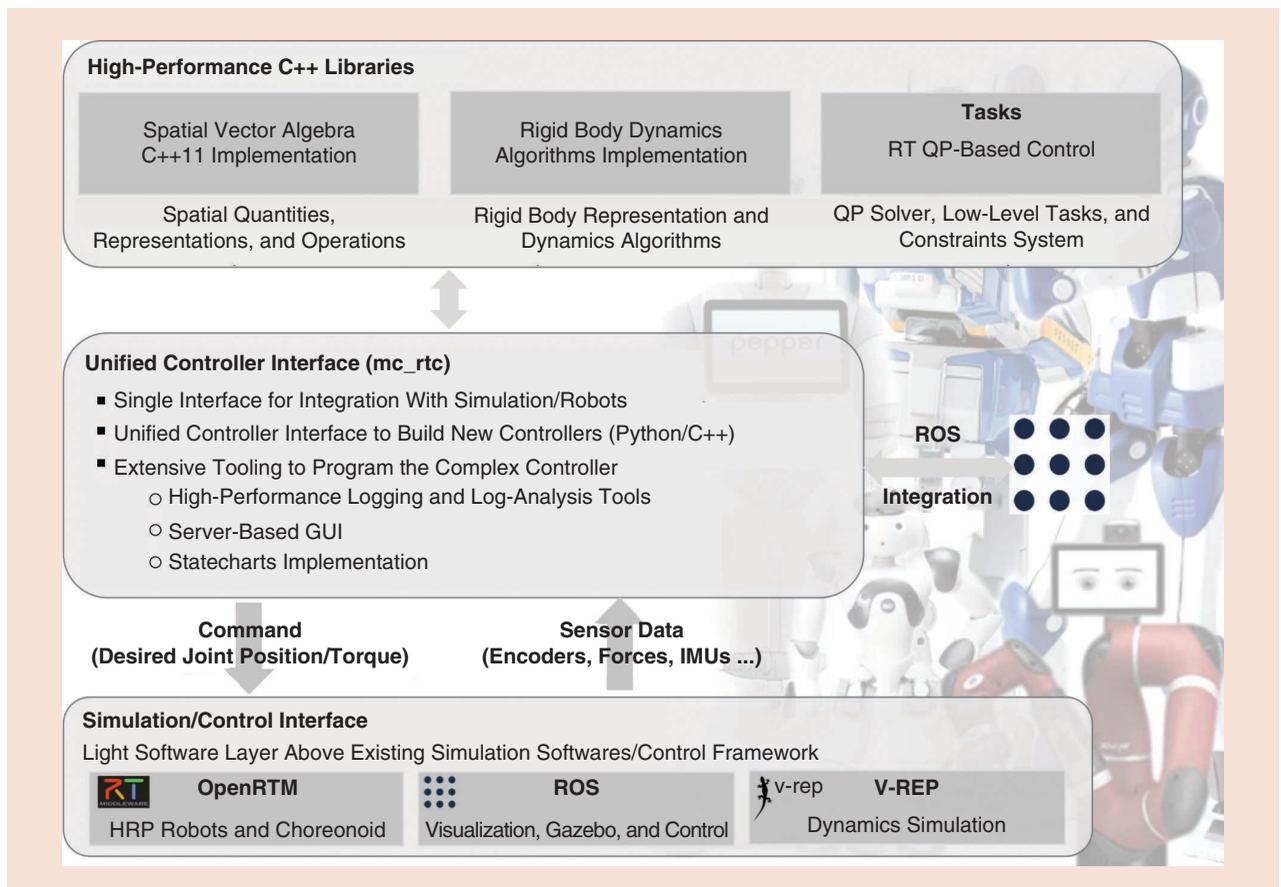


Figure 7. The control architecture of HRP-4. GUI: graphical user interface; IMU: inertial measurement unit.

(which are rather task-centric instead of model-centric), and a range of tools that support the development of new controllers (i.e., a new mission instantiation). This framework enables experts and nonexperts to build experiments.

We implemented walking and stair-climbing stabilization using `mc_rtc`. Stabilization, i.e., balance control, aims to correct the deviation of the floating base of the biped from its reference. With the CoM driven only by contact wrenches, this component consists of two main steps: 1) computing the desired contact wrenches and then 2) realizing these wrenches by force control. For the former, we applied proportional-integral-derivative feedback control of the DCM, followed by a QP-based wrench distribution to each foot in contact. For the latter, we used a single effector admittance controller to regulate the center of pressure (CoP) under each foot, a multi-effector admittance controller to set the foot pressures in double support, and a CoM admittance controller. The combination of these tasks realizes an overall whole-body admittance control scheme. The resulting stabilizer proved effective for both walking and climbing the factory staircase, the step height of which is 18.5 cm (see [26] for further details).

For walking-pattern generation, we implemented MPC using the open-source `copra` library. Viability was enforced using a terminal DCM condition at the end of the receding horizon, making walking trajectories two-step capturable. Phase timings were decided based on the scenario, i.e., 1.4/0.2 s for stair climbing, 0.7/0.1 s for sagittal walking, and 0.8/0.2 s for lateral stepping. Footholds were generated by an external footstep planner connecting current foot contacts to a world target provided by SLAM or the human operator. HRP-4 was able to lift itself up on its left leg but not on its right one due to a mechanical limitation issue. As a result, we selected footsteps that used only the left leg for lifting phases [26].

All of the tasks were programmed using `mc_rtc` statecharts. States are first programmed in C++ or Python classes; then, further states can be defined by specializing the configuration of an existing state. For example, a generically programmed state allows the defining of a set of tasks to perform until some completion criteria are reached for each of the tasks. The default behavior of this state is to add no tasks; however, a state can be configured to add a task on a given end effector and then define other states based upon this one with different targets for the end effector. Transitions are programmed by specifying a state, its output, and the next state. Orthogonal states and nested state machines are simply implemented as generic states. The walking operations are programmed using statecharts; hence, walking and

manipulation phases can be seamlessly integrated into a single controller. Walking stabilization can be enforced during manipulation by using orthogonal states.

SLAM technology was used to walk to predetermined walking targets (e.g., stairs, the bracket table, and bracket assembly) with subcentimeter precision. A keyframe graph map of the mockup was first generated by teleoperating the robot to walk within its intended workspace. As it walked, the robot actively looked at areas that maximize pose-tracking performance (e.g., well-textured surfaces and complex geometry). To ensure the best accuracy, walking targets were saved locally as the pose of the camera w.r.t. its nearest keyframe. The map generated was then used online to localize the robot. First, relocalization was achieved by comparing live images from the camera with recorded keyframes, providing the initial pose of the robot camera. Its relative transformation w.r.t. the recorded targets can then be easily computed through the keyframe graph. Due to inaccuracies in the keyframe graph (arising from camera calibration, pose-tracking drift while building the map, and so on), the accuracy of this transformation depends on, among other factors, the distance between the robot and its target (typically within a 5–10-cm range). As the robot approached its target, fewer keyframes were considered, and its accuracy increased. The final accuracy near the workspace target depended only on the local tracking error to the nearest keyframes. This camera-to-target transformation was converted into relative walking targets by forward kinematics.

Because our footstep planner is not online, walking to a target was decomposed in subplans. First, the robot walked toward a waypoint close to its intended target. Then, local adjustment steps were generated and executed based on the camera-to-target error. This process was repeated until the error converged below a desired threshold (i.e., <0.5 cm for stairs). With this method, the robot successfully and robustly reached each of its

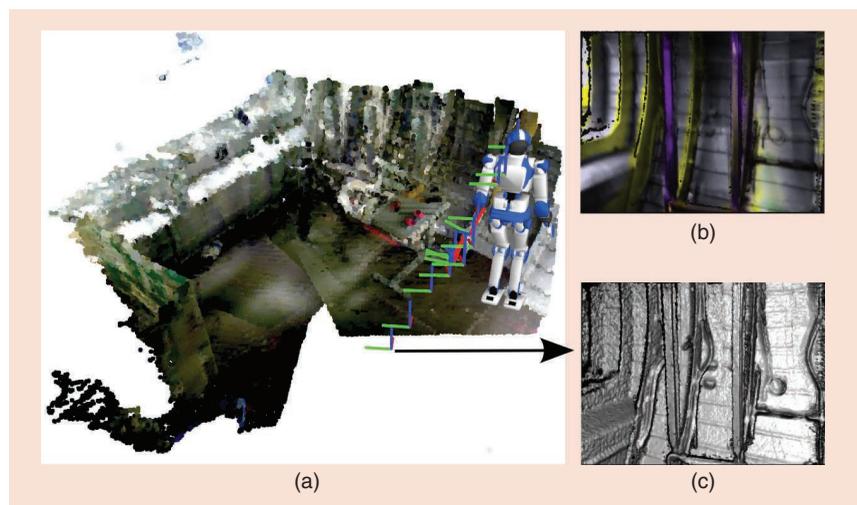


Figure 8. (a) A view of the 3D point cloud and the keyframe graph used by HRP-4 during the demo. (b) An image with semantic class overlay. (c) The fused 3D depth map.

workspace targets, starting from an initially unknown position, without the use of any specific or custom markers (see Figure 8).

The first task HRP-4 performed after climbing the stairs and entering the working area of the A350 mockup was grasping the bracket. This operation was programmed as follows (see Figure 9):

- position the left gripper above the bracket feeder
- trigger the tracking of the gripper and the feeder by ViSP [19], [20]
- trigger the visual servoing [21] to align the gripper above the bracket (3 cm), avoiding occlusions
- (without visual servoing) rotate the gripper, then lower it until a contact is sensed, then move slightly upward, and grasp and take the bracket
- prepare for locomotion toward the fuselage.

This process describes only the normal mode of operation (i.e., excluding failure recovery programmed in the statechart). For example, if the tracking of either the gripper or feeder is lost, the visual-servoing task is put on hold until tracking is recovered or a decision is made to resume the manipulation. The manipulation also involves two parallel behaviors, e.g., the stabilization described previously and gaze control, which ensures that HRP-4 continuously looks at the gripper during operation.

After grasping the bracket, HRP-4 walks toward the fuselage using SLAM to perform the bracket-assembly task, which is programmed as follows:

- 1) trigger environmental detection and tracking using ViSP
- 2) create contact between the right gripper and the A350 structure beam and keep a minimum pressure of 5 N

- 3) using ViSP, control the gripper to reach the fuselage at the bracket's desired location
- 4) glue the bracket, applying 10 N along the surface normally
- 5) resume the force to near zero, then release the bracket
- 6) remove the right-hand contact
- 7) step away from the frame and come to a safe position w.r.t. the tool table (using SLAM).

The final operation showcases the embedded, whole-body capacitive-sensing technology. We implemented a human-presence-awareness contactless compliance, which helps HRP-4 to avoid human closeness and prevent contact. The safety-repulsive field is defined from the detection range of the electrodes. A joint has a range of motion between an initial position and its lower or upper articular limit. The CoM task has a range between its initial position and a safe lower position along the world z -axis. In this demo, we achieved a simple electrode-to-joint or CoM-surjective mapping. The relation between an electrode capacitance measurement (ECM) value and the computed target for the desired task is expressed as a percentage in their respective range. Figure 10 illustrates the relationship between the sensor reading on the head electrodes (numbered 1 to 4) and the CoM height.

A configuration file describes a set of behaviors as {Name, Electrodes, Affect, and Target}, where *Name* is the behavior name, *Electrodes* is a list of electrodes, *Affect* is either a joint or task name (i.e., the HRP-4 joint name or CoM), and *Target* is either the lower or upper limit (i.e., an ECM equal to 0 means the task target is equal to the target value, and an ECM equal to maximum value means affect is in the initial position). Because explicitly describing every relation is too long (24 electrodes and seven targets for 13 behaviors), they

are compressed as follows: 1) head electrodes modify the CoM target along the world z -axis target; 2) shoulder electrodes modify the chest yaw joint target; 3) back electrodes modify the chest pitch joint target to lean forward, whereas the torso electrodes control backward leaning; and 4) the upper-arm electrodes affect shoulder roll and pitch.

Demonstrator of the TORO Humanoid Robot

TORO performed a complete demonstration for placing a bracket. In all phases of the presented demo, the robot was fully torque controlled. We used a passivity-based, whole-body controller [7], [8] for manipulating tasks and combined this controller with our DCM framework for walking [11]. The whole demonstration was executed in a fully autonomous way without any intervention by a human operator.

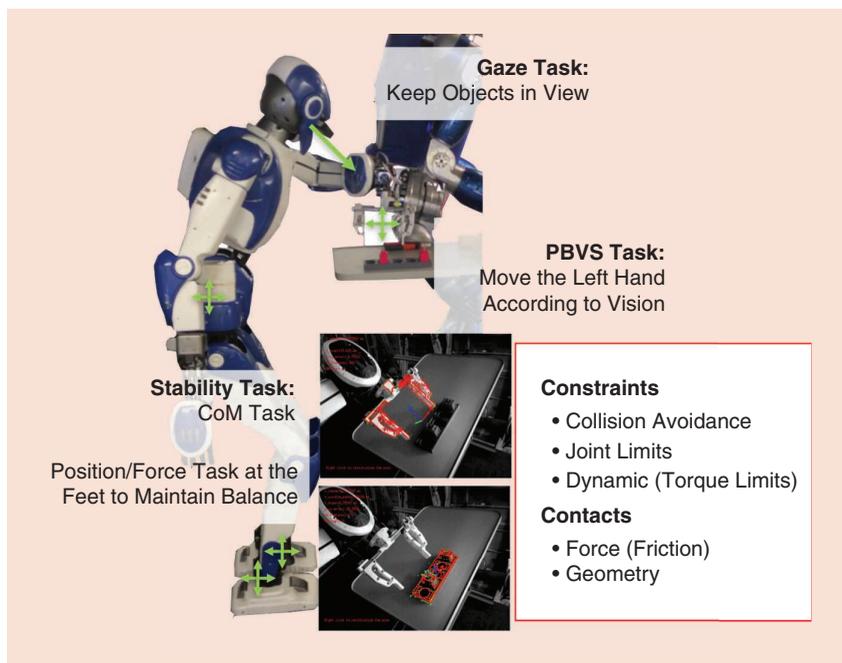


Figure 9. The QP tasks used during HRP-4's bracket grasping. PBVS: point-based visual servoing.

To navigate in the Airbus mockup, the robot uses landmarks, and the motion between two different landmarks defines a substage of step planning. Once the subgoal for a walking motion is achieved, the next stage is planned and executed. This approach tolerates the verification of self-localization of the robot within the mockup, to guarantee that the robot achieves the posture required to successfully fulfill the intended task, with an overall deviation lower than 2 cm w.r.t. the required position for placing the bracket. The walking motion was executed in different postures:

- normal, upright walking with knees slightly bent
- crouched walking, which is required to move within the mockup and prevent collisions with the head due to the ceiling's low height.

Based on the DCM formulation, CoM trajectories (i.e., position, velocity, and acceleration) were generated using a piecewise interpolation over a sequence of waypoints [27], while the trajectories of the feet were generated as fifth-order polynomials. These reference trajectories were tracked by a passivity-based whole-body controller.

Recent developments have permitted TORO to walk upstairs using toe-off motions that increase the kinematic workspace of the leg and therefore allow the robot to negotiate stairs with a step height and step length of 18 and 28 cm, respectively [28]. However, the demonstrator did not include entering and exiting the area using the stairs.

Figure 11 gives an overview of the complete operation. In this section, we briefly describe each of the individual steps.

- *Points 1 to 2:* First, the robot localizes itself w.r.t. the table on which the brackets are stored. Localization is performed using AprilTags. The robot approaches the table and moves into a configuration such that the bracket holder is well within the dexterous workspace of the left arm.
- *Point 3:* TORO finds the bracket holder's exact position on the table and picks up one of the brackets using a predefined grasp.
- *Point 4:* The robot backs away from the table and turns toward the direction of the mockup. Due to the ceiling's low height, the robot cannot walk fully upright within the mockup. It therefore changes its whole-body posture into a crouched configuration before entering the aircraft. In this configuration, the robot lowers its CoM by an additional 7 cm and bends its upper body forward 20°, resulting in a total

height of approximately 1.6 m (the normal walking height of the robot is 1.75 m).

- *Points 4 to 6:* The robot then walks toward tag 3 in the mockup and localizes itself. When it is close enough, it turns right

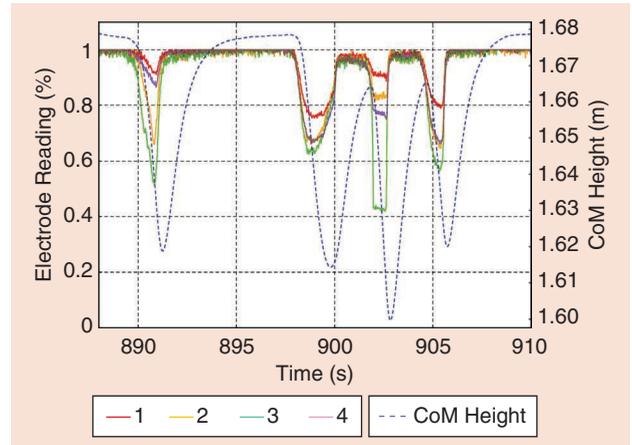


Figure 10. The interaction between the electrode and the CoM height shows how the robot reacts preemptively to contact. The user did not physically touch the robot (the electrode values approach, but do not hit, zero).

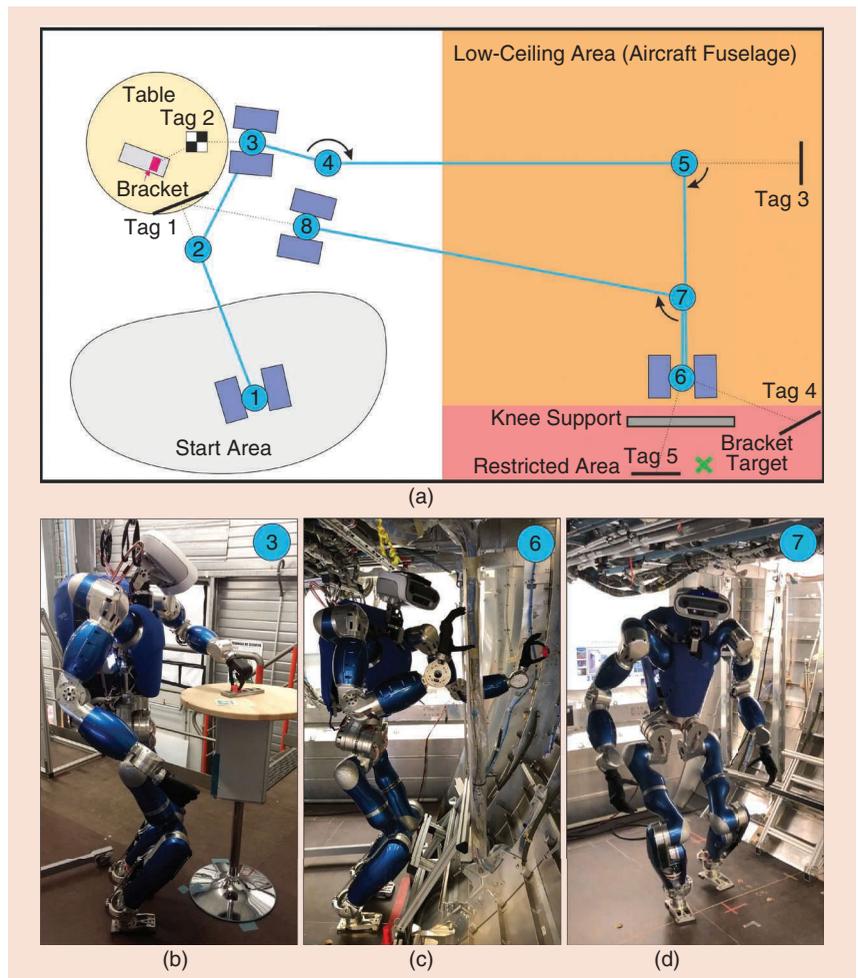


Figure 11. (a) An overview of the demonstration performed by TORO and photographs of TORO at (b) point 3, (c) point 6, and (d) point 7 in (a).

and localizes the final target location for the bracket. It further approaches the target location to a point from which it can place the bracket in a multicontact configuration.

- *Point 6:* The operation of placing the bracket can be achieved using three different multicontact configurations (see Figure 2):

- 1) A standing posture, with one arm grasping a vertical strut. This is a more suitable posture for executing the task, and the robot uses the additional support on the hand to extend its reachability while guaranteeing its balance.
- 2) A standing posture with knee support, where both knees rest on a horizontal beam. This posture is particularly interesting because it requires precise control of the internal forces in the robot. This is an over-constrained configuration because the contact forces at the feet and at the knees cannot be controlled fully independently.
- 3) A kneeling posture (i.e., no foot contact).

All of the postures exploit the combination of a passivity-based controller with a hierarchical whole-body control [8], which tolerates prioritization of the multiple tasks that the robot should execute, including support generation as its first priority and the CoM plus manipulation task as its second priority. In the presented demo, the robot utilizes the standing posture with knee support. In this way, the robot can lean forward to place the bracket in a location that would not be reachable using only feet support. The passivity-based balancer developed for TORO distinguishes between the end effectors used for balancing

(assuming that they are in rigid contact with the environment) and the end effectors used for accomplishing additional interaction tasks. A compliant reaction force acting on the CoM is then distributed to the end effectors used for balancing, while taking the desired control forces of the other end effectors into account. The algorithm requires the solution of an optimization problem, taking into account constraints for the unilaterality, friction, and position of the CoP [7], [9].

- *Points 7 to 8:* After the bracket is successfully placed, the robot retreats several steps, looks for the table holding the brackets (which is used for self-localization to find its way out of the mockup), and starts walking toward it.
- *Point 8:* Finally, the robot returns from the mockup and switches back to the upright posture. Then, the robot can restart the whole procedure to place the next bracket.

To demonstrate the robustness of torque-based walking using the DCM framework, we additionally demonstrate a short walk over uneven terrain (see Figure 12). This simulates possible ground irregularities such as cover plates, cables, or other debris lying on the floor.

Figure 13 illustrates TORO's control architecture. All of the control approaches used in this demonstration benefit from the joint torque control available in the robot. Although the low-level joint control runs at a rate of 3 kHz in local processors, the whole-body controllers run in a centralized, RT computer at a rate of 1 kHz. We used Linux with the PREEMPT_RT patch as the RT operating system. Additional details about TORO's hardware and software can be found in [29].

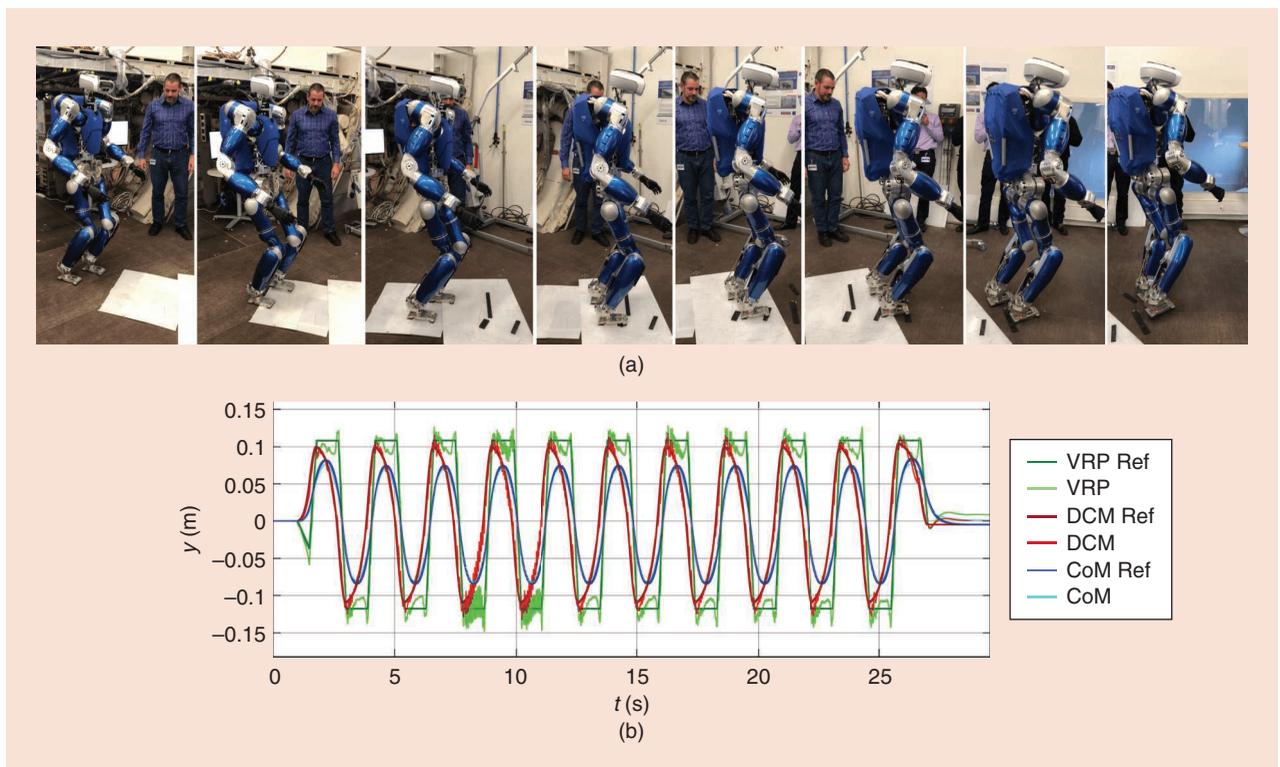


Figure 12. (a) The combined torque-based, whole-body control framework and DCM-based gait stabilization allow for walking over rough terrain. (b) The CoM, DCM, and virtual repealing point (VRP) ref trajectories. The increased oscillations of approximately $t = 10$ s are related to the edge contacts on the debris.

Discussion and Recommendations

The previously described experiment was prepared in situ during the course of two weeks. On the date of the final demonstration (corresponding to the final evaluation of the project), the experiment was performed two times with secured ropes. The reliability of the performance enabled the running of two other demonstrations without security harness/ropes. We aimed for the last one with HRP-4, but, unfortunately, a software bug caused it to fall on its right side. Some right arm parts of the cover appeared damaged, and the robot was fully repaired in lab. Because the mechanical piece linking the forearm to the arm was weak, it became deformed, and most of the shock was absorbed by it (thankfully, no actuators, gears, or capacitive sensing were damaged). This raises an interesting observation for designing more robust hardware.

Thus far, the primary limiting factor of our researches is the weakness of our platforms. Securing the humanoid robots acts as a high damping to our R&D efforts; therefore, we advise that any upcoming humanoid platform be designed, from the very beginning, to fall repeatedly without hardware damage. Working with the previously described humanoids, without any perspective of hardware improvement in the course of the project, put many constraints on what could be potentially achieved. We designed a new gripper for the HRP-4's grasping bracket and buttons, but that gripper could not be used for contact supporting robot motions. TORO uses a prosthetics hand as a gripper.

Nearly all of the limitations are mainly due to 1) nonproprietary technologies, 2) limited robustness in the perception, and 3) no recovery strategies being implemented because of subsequent engineering efforts. The consortium struggled with adapting perception, control, and planning to overcome these limitations. Full assessments and improving performance require extensive testing, which is difficult to achieve with humanoids. At this point in time, bridging the theory to practice is clearly another theory, i.e., there are still too many "magic numbers" to set, so it is worthwhile to research lowering the ad hoc tunings.

From the tests we conducted, range of motion (limited because of the wiring) is an important issue in cluttered/confined environments. In fact, extra redundancy, including the inconvenience of additional costs, is also needed in some places. All of our robots have hard covers, which are not robust to secure contact formation. Intrinsic or active joint compliance is not enough to secure a sustained contact (as the structure of the airplane is also rigid), whereas soft covers are a better fit and offer many other advantages. Safety is still an open issue; it is worth dedicating efforts to extend the cobots' safety requirements International Standards Organization (ISO) 10218 standard parts 1 and 2 and ISO/technical specification 15066 to include humanoid robots. As for manipulation capabilities, we may consider interchangeable tools instead of a single sophisticated gripper. Concerning the energy and power supply, Airbus

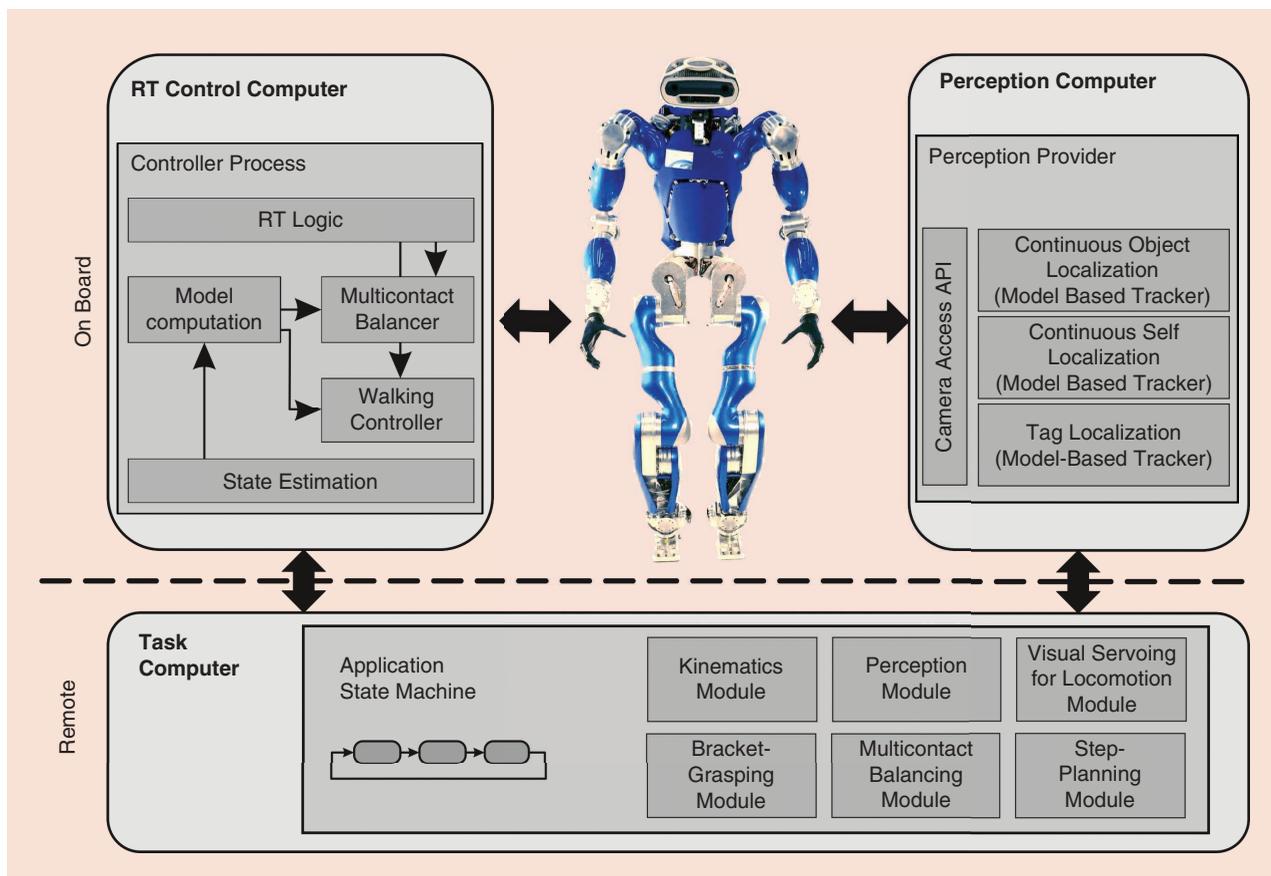


Figure 13. The control architecture of TORO. API: application programming interface.

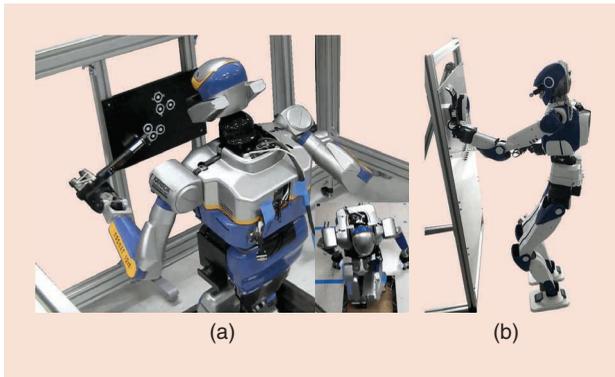


Figure 14. (a) The HRP-2Kaï's narrow access in multicontact and fast, dynamic torqueing enables the handling of four contacts. (b) Circuit breaker checking combined with HRP-4's multisensory task-space control is achieved without visual markers.

allows power supply cables nearly everywhere; therefore, the humanoid robots can be embedded with short-time autonomy as long as they have the capability to autonomously plug and unplug themselves to the available power supply cables. There is still much to accomplish through the use of impact absorption technology for safe falling as well as in falling detection and recovery strategies. There are also many other recommendations that we will use for the innovation stage.

There are encouraging results found in the software architecture (i.e., the HRP-4 is fully open source code). We were able to achieve (using the exact same controller, task templates, and methodology) two other use cases that we did not report on in this article because they were part of another project and were not demonstrated in situ: i.e., torqueing in multicontact dynamic motion and circuit breaker operations [3] but without the markers (see Figure 14). This is very promising, as our new objective is to reach human-speed performance and reliability in all of these tasks and future ones as well.

Any further continuation of this project is subject to the ultimate question by Airbus as well as by any other large-scale manufacturer that shows interest in this technology (and there are already a few): Who will build and commercialize humanoid robots for manufacturing—that is to say, certified robots that are robust enough to work continuously with a near-human-speed performance, including a time cycle and reliability comparable to current industrial robots? The few big companies capable of doing so do not see significant opportunities w.r.t. their current/other products, whereas medium- or small-size companies are certainly not large enough for such a challenge, knowing that it does not suffice to provide the hardware. An integrated automation turnkey solution is also needed.

References

[1] K. Bouyarmane, S. Caron, A. Escande, and A. Kheddar, "Multi-contact planning and control," in *Humanoid Robotics: A Reference*, A. Goswami and P. Vadakkepat, Eds. Dordrecht, The Netherlands: Springer, 2019, pp. 1763–1804. doi: 10.1007/978-94-007-6046-2_32.

[2] S. Brossette, A. Escande, and A. Kheddar, "Multi-contact postures computation on manifolds," *IEEE Trans. Robot.*, vol. 34, no. 5, pp. 1252–1265, 2018.

[3] A. Bolotnikova et al., "A circuit-breaker use-case operated by a humanoid in aircraft manufacturing," in *Proc. IEEE Conf. Automation Science and Engineering*, 2017, pp. 15–22.

[4] G. Mesesan, J. Engelsberger, B. Henze, and C. Ott, "Dynamic multi-contact transitions for humanoid robots using divergent component of motion," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2017, pp. 4108–4115.

[5] K. Bouyarmane, K. Chappellet, J. Vaillant, and A. Kheddar, "Quadratic programming for multirobot and task-space force control," *IEEE Trans. Robot.*, vol. 35, no. 1, pp. 64–77, 2019.

[6] H. Audren and A. Kheddar, "3D robust stability polyhedron in multi-contact," *IEEE Trans. Robot.*, vol. 34, no. 2, pp. 388–403, 2018.

[7] B. Henze, M. A. Roa, and C. Ott, "Passivity-based whole-body balancing for torque-controlled humanoid robots in multi-contact scenarios," *Int. J. Robot. Res.*, vol. 35, no. 12, pp. 1522–1543, 2016.

[8] B. Henze, A. Dietrich, M. A. Roa, and C. Ott, "Multi-contact balancing of humanoid robots in confined spaces: Utilizing knee contacts," in *Proc. IEEE Int. Conf. Intelligent Robots and Systems*, 2017, pp. 679–704.

[9] C. Ott and S.-H. Hyon, "Torque-based balancing," in *Humanoid Robotics: A Reference*, A. Goswami and P. Vadakkepat, Eds. Dordrecht, The Netherlands: Springer, 2016, pp. 1–26. doi: 10.1007/978-94-007-7194-9_39-1.

[10] P.-B. Wieber, R. Tedrake, and S. Kuindersma, "Modeling and control of legged systems," in *Springer Handbook of Robotics*, 2nd ed., B. Siciliano and O. Khatib, Eds. New York: Springer-Verlag, 2016, pp. 1203–1234.

[11] J. Engelsberger, C. Ott, and A. Albu-Schäffer, "Three-dimensional bipedal walking control based on divergent component of motion," *IEEE Trans. Robot.*, vol. 31, no. 2, pp. 355–368, 2015.

[12] N. Bohórquez, A. Sherikov, D. Dimitrov, and P.-B. Wieber, "Safe navigation strategies for a biped robot walking in a crowd," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, 2016, pp. 379–386.

[13] N. Scianca, M. Cognetti, D. De Simone, L. Lanari, and G. Oriolo, "Intrinsically stable MPC for humanoid gait generation," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, 2016, pp. 601–606.

[14] D. De Simone, N. Scianca, P. Ferrari, L. Lanari, and G. Oriolo, "MPC-based humanoid pursuit-evasion in the presence of obstacles," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2017, pp. 5245–5250.

[15] A. Pajon and P.-B. Wieber, "Safe 3D bipedal walking through linear MPC with 3D capturability," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2019, pp. 1404–1409. doi: 10.1109/ICRA.2019.8794117.

[16] M. Meilland and A. I. Comport, "On unifying key-frame and voxel-based dense visual slam at large scales," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2013, pp. 3677–3683.

[17] F. I. Ireta Munoz and A. I. Comport, "Point-to-hyperplane ICP: Fusing different metric measurements for pose estimation," *Adv. Robot.*, vol. 32, no. 4, pp. 161–175, 2018.

[18] H. Mahé, D. Marraud, and A. I. Comport, "Real-time RGB-D semantic keyframe slam based on image segmentation learning from industrial CAD models," submitted for publication.

[19] E. Marchand, F. Spindler, and F. Chaumette, "ViSP for visual servoing: A generic software platform with a wide class of robot control skills," *IEEE Robot. Autom. Mag.*, vol. 12, no. 4, pp. 40–52, 2005.

[20] S. Trinh, F. Spindler, E. Marchand, and F. Chaumette, "A modular framework for model-based visual tracking using edge, texture and

depth features,” in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2018, pp. 89–96.

[21] D. J. Agravante, G. Claudio, F. Spindler, and F. Chaumette, “Visual servoing in an optimization framework for the whole-body control of humanoid robots,” *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 608–615, 2017.

[22] N. Scianca, L. Lanari, and G. Oriolo, “Deliverable D3.3: Final safety guidelines and strategies,” COMANOID, Paris, 2019. [Online]. Available: <http://comanoid.cnrs.fr/deliverables>

[23] V. Samy and A. Kheddar, “Falls control using posture reshaping and active compliance,” in *Proc. IEEE-RAS 15th Int. Conf. Humanoid Robots*, 2015, pp. 908–913.

[24] T. Mattioli and M. Vendittelli, “Interaction force reconstruction for humanoid robots,” *IEEE Robot. Autom. Lett.*, vol. 2, no. 1, pp. 282–289, 2017.

[25] G. Garofalo, N. Mansfeld, J. Jankowski, and C. Ott, “Sliding mode momentum observers for estimation of external torques and joint acceleration,” in *Proc. IEEE Int. Conf. Robotics and Automation*, 2019. doi: 10.1109/ICRA.2019.8793529.

[26] S. Caron, A. Kheddar, and O. Tempier, “Stair climbing stabilization of the HRP-4 humanoid robot using whole-body admittance control,” in *Proc. Int. Conf. Robotics and Automation*, 2019. doi: 10.1109/ICRA.2019.8794348.

[27] G. Mesesan, J. Engelsberger, C. Ott, and A. Albu-Schäffer, “Convex properties of center-of-mass trajectories for locomotion based on divergent component of motion,” *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3449–3456, 2018.

[28] B. Henze, M. A. Roa, A. Werner, A. Dietrich, C. Ott, and A. Albu-Schäffer, “Analysis of human-like behaviors in a humanoid robot: Quasi-static balancing using toe-off motion and stretched knees,” in *Proc. IEEE Int. Conf. Robotics and Automation*, 2019. doi: 10.1109/ICRA.2019.8794096.

[29] C. Ott et al., “Mechanisms and design of DLR humanoid robots,” in *Humanoid Robotics: A Reference*, A. Goswami and P. Vadakkepat, Eds. Dordrecht, The Netherlands: Springer, 2016, pp. 1–26. doi: 10.1007/978-94-007-7194-9_132-1.

[30] Comanoid. [Online]. Available: comanoid.cnrs.fr

[31] Fogale Robotics, “Sensitive surfaces for human/robot interactions & cooperation.” [Online]. Available: www.fogale-robotics.com

Abderrahmane Kheddar, National Center for Scientific Research, University of Montpellier, Laboratory for Computer Science, Robotics and Microelectronics of Montpellier, France. Email: kheddar@lirmm.fr.

Stéphane Caron, National Center for Scientific Research, University of Montpellier, Laboratory for Computer Science, Robotics and Microelectronics of Montpellier, France. Email: stephane.caron@lirmm.fr.

Pierre Gergondet, National Center for Scientific Research, National Institute of Advanced Industrial Science and Technology, Joint Robotics Laboratory, Tsukuba, Japan. Email: pierre.gergondet@gmail.com.

Andrew Comport, National Center for Scientific Research, University de Nice, France. Email: andrew.comport@cnrs.fr.

Arnaud Tanguy, University of Montpellier, Laboratory for Computer Science, Robotics and Microelectronics of Montpellier, France. Email: arnaud.tanguy@lirmm.fr.

Christian Ott, German Aerospace Center, Oberpfaffenhofen, Germany. Email: christian.ott@dlr.de.

Bernd Henze, German Aerospace Center, Oberpfaffenhofen, Germany. Email: bernd.henze@dlr.de.

George Mesesan, German Aerospace Center, Oberpfaffenhofen, Germany. Email: george.mesanan@dlr.de.

Johannes Engelsberger, German Aerospace Center, Oberpfaffenhofen, Germany. Email: johannes.engelsberger@dlr.de.

Máximo A. Roa, German Aerospace Center, Oberpfaffenhofen, Germany. Email: maximo.roa@dlr.de.

Pierre-Brice Wieber, University Grenoble Alpes, Inria, France. Email: pierre-brice.wieber@inria.fr.

François Chaumette, National Institute for Research in Computer Science and Automation, Irisa Rennes, France. Email: francois.chaumette@irisa.fr.

Fabien Spindler, National Institute for Research in Computer Science and Automation, Irisa Rennes, France. Email: fabien.spindler@irisa.fr.

Giuseppe Oriolo, Sapienza University of Rome, Italy. Email: oriolo@diag.uniroma1.it.

Leonardo Lanari, Sapienza University of Rome, Italy. Email: lanari@diag.uniroma1.it.

Adrien Escande, National Center for Scientific Research, National Institute of Advanced Industrial Science and Technology, Joint Robotics Laboratory, Tsukuba, Japan. Email: adrien.escande@aist.go.jp.

Kevin Chappellet, National Center for Scientific Research, National Institute of Advanced Industrial Science and Technology, Joint Robotics Laboratory, Tsukuba, Japan. Email: chappellet.kevin@gmail.com.

Fumio Kanehiro, National Center for Scientific Research, National Institute of Advanced Industrial Science and Technology, Joint Robotics Laboratory, Tsukuba, Japan. Email: f-kanehiro@aist.go.jp.

Patrice Rabaté, Airbus, Saint-Nazaire, France. Email: patrice.rabate@airbus.com.

