

Center-of-Mass-Based Grasp Pose Adaptation Using 3D Range and Force/Torque Sensing

Dimitrios Kanoulas*, Jinoh Lee, Darwin G. Caldwell
and Nikos G. Tsagarakis

*Humanoids and Human-Centered Mechatronics Lab,
Istituto Italiano di Tecnologia (IIT)
Via Morego 30, 16163 Genova, Italy
Dimitrios.Kanoulas@iit.it

Received 26 August 2016
Accepted 8 January 2018
Published

Lifting objects, whose mass may produce high wrist torques that exceed the hardware strength limits, could lead to unstable grasps or serious robot damage. This work introduces a new Center-of-Mass (CoM)-based grasp pose adaptation method, for picking up objects using a combination of exteroceptive 3D perception and proprioceptive force/torque sensor feedback. The method works in two iterative stages to provide reliable and wrist torque efficient grasps. Initially, a geometric object CoM is estimated from the input range data. In the first stage, a set of hand-size handle grasps are localized on the object and the closest to its CoM is selected for grasping. In the second stage, the object is lifted using a single arm, while the force and torque readings from the sensor on the wrist are monitored. Based on these readings, a displacement to the new CoM estimation is calculated. The object is released and the process is repeated until the wrist torque effort is minimized. The advantage of our method is the blending of both exteroceptive (3D range) and proprioceptive (force/torque) sensing for finding the grasp location that minimizes the wrist effort, potentially improving the reliability of the grasping and the subsequent manipulation task. We experimentally validate the proposed method by executing a number of tests on a set of objects that include handles, using the humanoid robot WALK-MAN.

Keywords: 3D Perception; range sensing; force/torque sensing; grasp adaptation; manipulation; object contacts; exteroceptive sensing; proprioceptive sensing; humanoids.

1. Introduction

Grasping strategies for object manipulation have been extensively studied over the past few years,^{43,8,36,32} especially for picking up light weight objects. Recent advancement in control, planning, and perception have enabled robots to complete various manipulation tasks mostly considering the geometry of the object shape and largely excluding the weight and/or the mass distribution of the grasped/manipulated object. Therefore, grasping, picking up, and eventually carrying heavy objects is considered an open problem and very challenging, particu-

larly when the distribution of mass is not exteroceptively detectable. Perception is a main aspect for completing these tasks, especially when grasp reliability during manipulation and efficient robot joints load reduction are required, while torque limitations also exist for the robotic arms. In recent works, grasp detection has been usually achieved using either exteroceptive perception, such as 2D/3D visual⁶⁵ or tactile^{70,22} sensing, where the mass distribution is not considered, or proprioceptive perception, such as force/torque sensing³¹ or the robot kinematics for contact detection. Our approach uses a combination of exteroceptive and proprioceptive perception to improve grasping.

Inspired from the debris task during the DARPA Robotics Challenge 2015, where wood debris pieces had to be removed from the robot's path, we present a new grasping method for a similar type of objects (Fig. 1). The challenge when only exteroceptive perception is used is that the object weight and mass distribution are unknown, leading often to grasps that may generate high wrist moments. These can eventually result in object drops (especially for underactuated hands) or robot overloading and instability. The grasping force and wrist torque become the bottleneck in these scenarios, since there are strict hardware strength force/torque limits. To improve the reliability of such a grasp during manipulation, we present a novel method that combines 3D range and wrist force/torque sensing to detect the Center-of-Mass (CoM)-based grasp pose for objects which include handles and lie on support surfaces, like a tabletop. In this paper, we study the case of single arm grasps, considering whole-body balancing.

The key aspect of our approach is the human-inspired hypothesis that holding heavy objects closer to their CoM makes the grasping more reliable and decreases the wrist torque effort when lifting them. In addition, the risk of reaching the wrist

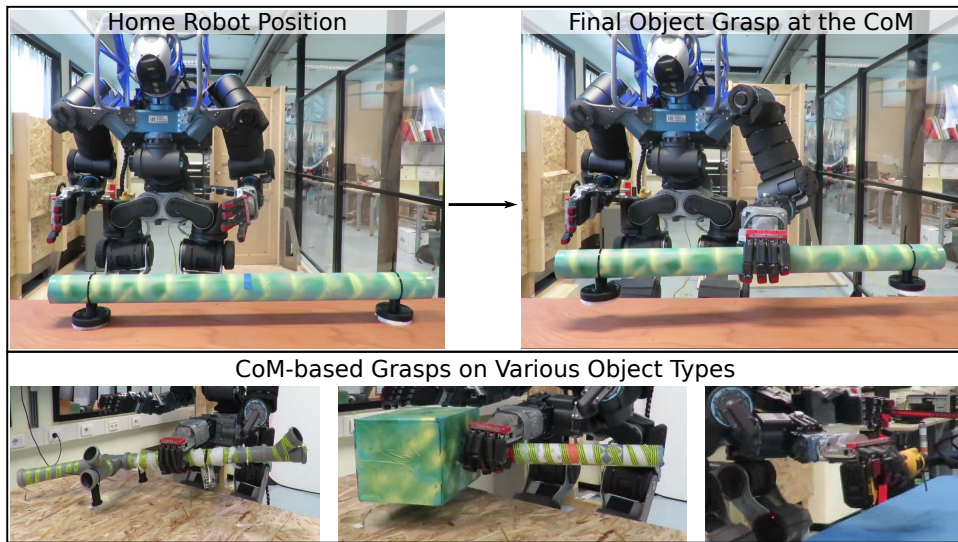


Fig. 1: The WALK-MAN humanoid robot and the proposed grasping for four different object types.

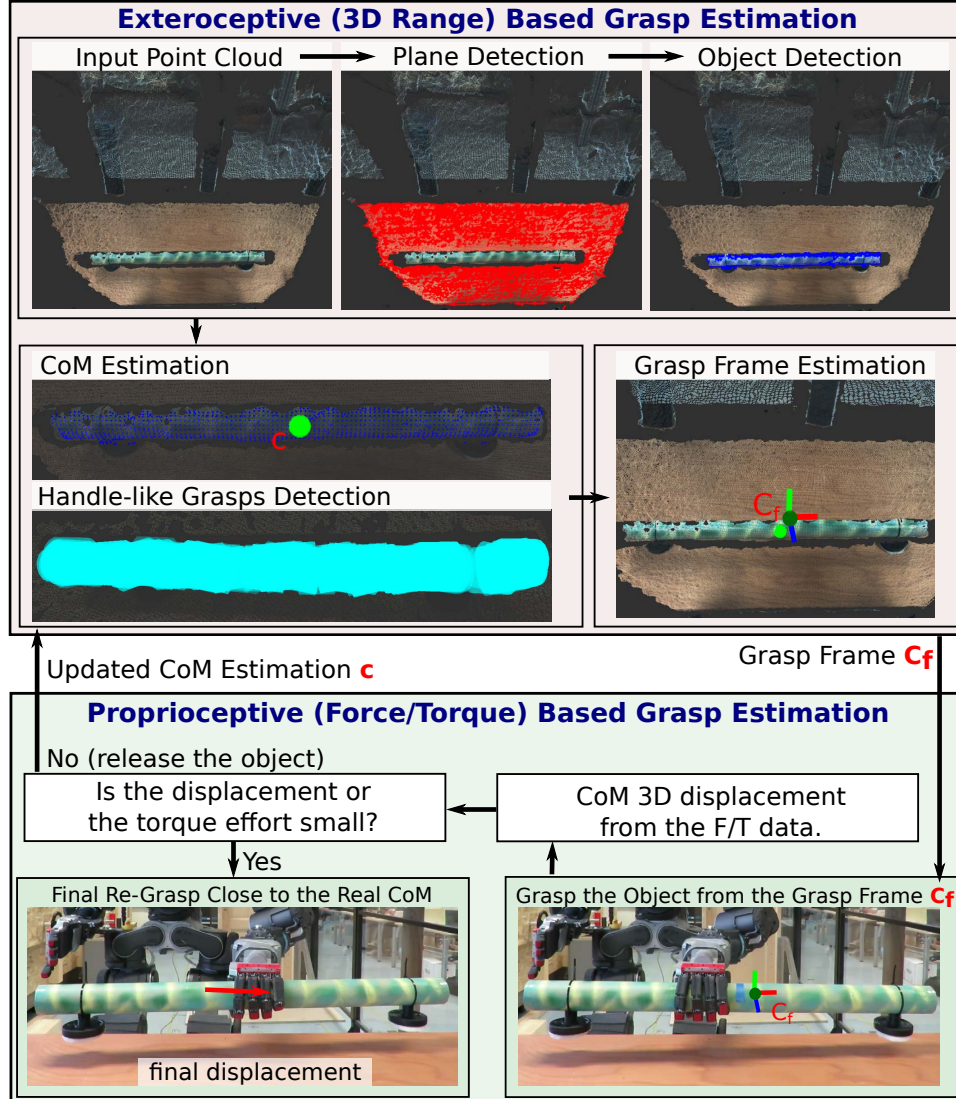


Fig. 2: System overview showing the main steps of the CoM-based grasp adaptation algorithm.

torque limits is potentially reduced. Our method is divided into two iterative stages (Fig. 2). Initially, the object's CoM is estimated geometrically using voxelized 3D range data, under the assumption that it is made by isotropic material with constant mass density. In the first exteroceptive-based stage, a set of handle-like (cylindrical) grasps are localized on the object. From them, the closest to the estimated CoM is selected to grasp, minimizing the wrist torque effort among the several handle options (Sec. 3.1). In the second proprioceptive-based stage, the robot lifts the object using a single arm, measuring in the meantime the wrist force/torque data.

A new CoM is estimated using these data, leading back to the first stage until a minimum wrist torque has been reached (Sec. 3.2). It is worth noting that in our method, we achieve grasps, even when the actual CoM of the object does not lie in the object itself or inside a handle area.

Next, we review the research context, followed by a review of the handle-like grasp representation, as well as the robotic platform description (Sec. 2). We then present in detail the 3D range exteroceptive perception system (Sec. 3.1) and the force/torque proprioceptive one (Sec. 3.2). Finally, we present experimental results using the WALK-MAN humanoid robot⁶⁶ on grasping heavy objects that include handles (Sec. 4). The system is implemented in C++ using the Point Cloud⁵⁹ and the Surface Patch^{28,26} Library, and is part of the source code designed for the DARPA Robotics Challenge Finals in 2015, used from the WALK-MAN team.

1.1. *Related work*

Object manipulation is considered one of the basic fields in robotics.⁴⁸ Both exteroceptive and proprioceptive perception were used to improve task completions. Grasping involves either object picking/moving or tool manipulation tasks, where a tool needs to be grasped from a particular position to be used accordingly.

The area of picking up objects using perception was mainly focused on lightweight objects, considering mostly the geometry of the object shape and excluding the weight or mass distribution from the selection of the grasping point. Recently, range sensing was extensively used for localizing grasps. For instance, range data acquired from low cost structured light sensors were used to extract geometrically meaningful grasps such as cylindrical handles,⁶⁴ curved patches,²⁷ or antipodal points⁶⁵ on light toy/kitchen objects to complete an empty-the-basket task. Similarly, learning approaches on RGB-D data were used to improve grasping while clearing piles of toy objects,^{14,12,15} using also geometric representations such as rectangles.²³ More recently, deep learning approaches were also considered in the literature, increasing the grasping success rate for similar type of objects,^{37,16,57,38} mainly applied on the Baxter robot which uses grippers. Interactive approaches^{29,11} were also used, where range data were checked for changes to verify and improve starting grasp hypotheses. Geometric feature and template matching in 3D point clouds^{33,18,63} were also developed for grasp selection and planning. All these methods were designed to work with high success rates for novel objects in clutter, using only range sensing. The visual sensing limitations for heavier objects is been considered in our work, which defers from all the above ones.

Searching for tool object affordances in range and RGB images for a particular use, was also considered in the literature.^{61,67,56,24,25} Additionally, interaction between the robot and the environment also played a role in self-learning affordances.^{46,42} More recently, deep learning was also used for pixelwise affordances classification.^{49,50,52} Usually, tools such as drills and hammers need to be grasped in a particular way, which make the problem different than what we study in this

paper.

Other types of sensing were also used to complete manipulation tasks. For instance, tactile sensing^{55,40,53} was used for small object localization and manipulation (see Ref. ⁷⁰ for a review). Force/torque sensing was mainly used in the literature for the detection of contact,³¹ slippage,⁶⁹ or shape² prediction. An alternative approach to more stable grasping is through the object model learning. For instance, tracking contacts while estimating the object’s dimensions, mass, and friction⁷¹ or updating the object’s attribute using tactile sensing³⁵ was studied in a probabilistic framework. Towards active manipulation,⁵⁴ object six Degrees-of-Freedom (DoF) localization takes place using various methods, such as the Scaling Series via touching¹ or using Bayesian approaches for tactile sensing.⁶⁸ In a much different direction, learning methods were also used for whole-body manipulation, for instance to learn friction models of the objects.⁶²

Small object localization for in-hand manipulation was also extensively studied, using hybrid sensing methods. In particular, vision, force, and tactile sensing was used in various sensing fusion combinations to compute both finger contacts and the applied forces during grasping.³ Stereo vision and wrist force/torque sensing was used in combination with joint position sensing, to localize fingers with respect to object faces,¹⁷ while tactile and vision sensing was used to localize objects, providing robustness to occlusions and sensor failures for multi-fingered hands, in static configurations²¹ or during manipulation.⁹

Task-oriented grasping methods using vision, proprioception, and tactile sensing⁷ to increase stability were introduced in the direction of tool use. Grasp adaptation in the controller level (stiffness) was also introduced³⁹ to increase the performance, while regrasps were also used to reorient small objects in the environment.⁵¹ Grasp planning was studied, such that grasps can guarantee lack of slippage and resistance to perturbations, using geometric object models and their theoretical CoM position and inertia.⁴¹ CoM estimation from wrist force/torque data without the use of vision has an early history in the literature.⁴ In the closest work to our method,^{6,34} learning techniques for tactile coupled with vision sensing were developed to make a grasp more stable. Still regrasping was not studied in any of these works. To our knowledge, we are the first to use 3D range and wrist force/torque sensing iteratively to regrasp heavy objects with irregular mass distribution, based on the estimated CoM position for more stable and torque efficient manipulation.

2. Grasp Representation and the Robotic System

The goal of our method is the detection of a reliable and torque efficient grasp. Such a grasp pose should be as close as possible to the object’s CoM, where the applied torque is the minimum. The final grasp localization is estimated in two iterative stages until a termination threshold criterion is met, using exteroceptive and proprioceptive perception. Note that for simplicity and clarity, in the next two

sections, we present some of the results on a simple cylindrical object, but later in the experimental section, we show that the method is generic to any object that includes handles.

Representing and localizing grasps in range data is well studied in robotic manipulation. We are interested in the problem of finding all the possible graspable areas on the object. For this reason, we apply one of the state-of-the-art methods⁶⁴ that uses cylindrical handles to geometrically represent grasp affordances. The original paper focuses on lifting light objects by grasping the closest handle on the object. In this paper, we extend the idea by a more sophisticated CoM-based grasp selection. First, we briefly review the handle-like grasp affordances representation and localization algorithm using 3D point clouds. We then present the robotic platform including both the exteroceptive and proprioceptive sensing system and the robotic hand that has been used.

2.1. *Grasp representation and localization*

A grasp handle (illustrated in cyan color in the figures) is represented as a cylindrical shell, which is a set of a fixed number of co-linear cylinders of different radii. Each cylinder is parametrized by its centroid, major principal axis, and radius.

The localization algorithm works as follows: Initially a set of uniform 3D points are sampled from the cloud that was acquired from a range sensor. For each of these points, a local spherical point cloud neighborhood of fixed size (between 2cm and 3cm) is extracted. Then, a quadric surface is algebraically fitted to the neighborhood, using Taubin’s normalization method. Given the curvatures along the two principal axis coming from the fitting, only those neighborhoods below some parametrized thresholds are considered. For each one of those, a cylinder is fitted, assuring non-collision with surfaces during the object manipulation (i.e., a gap without points around the cylinder). Handles are fixed sets of co-linear cylinders that are checked against some manually parametrized thresholds of their centroids, principal axes, and radii distances that form the final enveloping grasp affordance set. More details can be found in the original paper.⁶⁴ In this work, as described in Sec. 3.1, we extract a set of handles on the object of interest that samples it uniformly random, but very densely.

2.2. *Robotic platform*

For the experiments, the WALK-MAN electrical motor driven humanoid robot has been used as shown in Fig. 3(a). WALK-MAN has 31 DoF, with two actuators for its hands, while it is 1.85m tall and weighs 118kg. Visual sensing includes the CMU Multisense-SL system, which has a stereo and a LiDAR sensor, while four 6 DoF force/torque sensors are attached in the two wrists and ankles. The hands are customized from the Pisa/IIT SoftHands¹⁰ of 11x11cm palm size and 12cm finger length. During a grasp, the hand frame H_f as appears in Fig. 3(b), needs to co-align

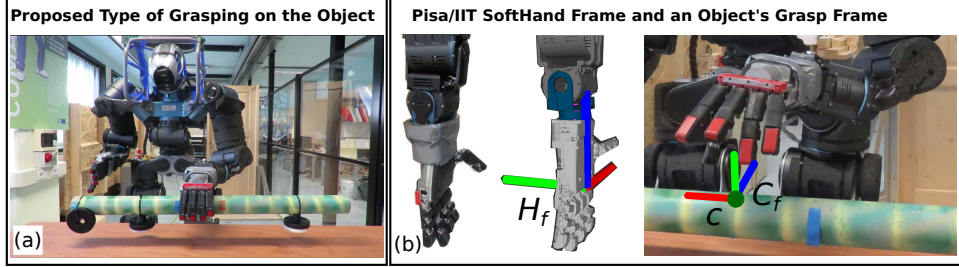


Fig. 3: (a) An object grasp with the WALK-MAN humanoid robot. (b) The real and the simulated Pisa/IIT SoftHand with its frame H_f and a grasp frame C_f with its origin point c on an object.

with the detected object's grasp frame C_f at the origin point c . This grasp frame will be the handle frame that is closest to the CoM estimation.

There are three big challenges for the manipulation tasks using the particular robotic platform. The first one has to do with the noisy stereo camera data, the second with the instabilities of the underactuated hand grasps, and the third with the robot's balancing during the execution of the task. These aspects are going to be considered in the following sections. It is worth noting that grasp stability is importantly benefited from the active Pisa/IIT hand and passive WALK-MAN robot joint compliance. Using the particular hardware helps with small grasp uncertainties either due to kinematic error accumulation or due to inaccurate grasp localization. Moreover, we benefit from the finger compliance that enables the hand to envelop the contact surface, avoiding complicated control schemes. When more dexterous manipulations are needed, for instance, manipulating smaller objects, the particular underactuated hand may be challenging to use.

3. CoM-Based Grasp Pose Adaptation Method

The introduced CoM-based grasp pose adaptation method includes two iterative stages (Fig. 2). In the first one (Sec. 3.1), 3D visual handle grasps are localized on the object and the closest to the object's CoM is selected. In the second stage (Sec. 3.2), the object is lifted and its CoM is estimated using the wrist force/torque data. The stages are repeated until the minimum wrist torque effort is reached.

3.1. Exteroceptive-based grasp estimation

To initially estimate the CoM, we use the 3D range point cloud data, acquired from the robot's range sensor. In CAD systems, the CoM is detected by splitting the object into voxels and averaging the weighted distance from a fixed point, where the weights represent the object's density at the particular cube.³⁰ Similarly, the proposed method *visually* detects the initial CoM position of the object, using 3D voxelization and handle-grasps localization over the input object data, with the following steps, called 'Stage I'. (see also Fig. 4):

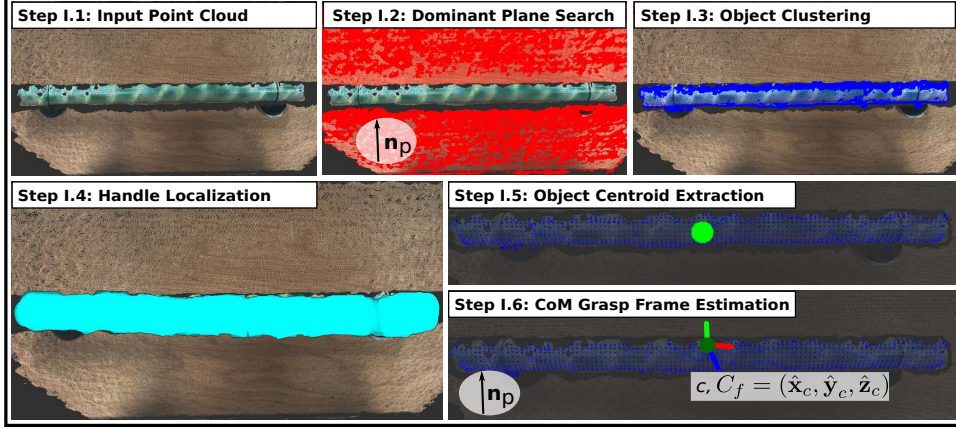


Fig. 4: The six steps of the 3D range grasp frame localization. In red is the dominant plane, in cyan, the handle-like (cylindrical) grasps, in light green, the estimated CoM in the voxelized cloud, and in dark green the refined selected grasp position \mathbf{c} with its frame C_f .

Stage I (3D range-based grasp localization)

- **Step I.1 [input point cloud]:** Acquire a point cloud from the range sensor.
- **Step I.2 [dominant plane search]:** Find the dominant plane (with its normal \mathbf{n}_p) and segment the points above it.
- **Step I.3 [object clustering]:** Cluster the segmented points into objects, using their Euclidean distances and then extract the largest one.
- **Step I.4 [handle localization]:** Fit a set of handle-like grasps of the size of the hand⁶⁴ to oversample the whole object cloud.
- **Step I.5 [object centroid extraction]:** Split the segmented object cloud space into fixed-size 3D voxels and for each one, compute its centroid; the median of the centroids represents the CoM position.
- **Step I.6 [CoM grasp frame estimation]:** From the extracted handle-like grasps, select the closest to the CoM estimation.

3.1.1. The algorithm

The CMU Multisense-SL stereo camera provides an organized 1024×1024 point cloud in 15Hz framerate. After acquiring a cloud, we first filter out those points that are not reachable, to speed up the following computations. These are the points that are approximately further than the distance that the hands can reach when they are in full extend, i.e., roughly 1.5m away from the camera frame, given that the lower body is not used in this paper.

Assuming that the object lies on a flat surface, e.g., a tabletop, RANdom SAMpling Consensus (RANSAC) clustering procedure is used to extract the dominant plane cloud P in the scene,¹⁹ with the angles between the point normals as the

classification criterion. The local normal vector for each point is computed using the integral images method²⁰ (i.e., covariance matrix estimation). The extracted plane’s normal vector towards the camera viewpoint will be denoted as \mathbf{n}_p .

To extract the points that are above the support surface in the direction of the normal vector \mathbf{n}_p , the plane cloud P is first projected onto the fitted plane, a convex hull of the projected table points is created and all the rest points are then projected on the same plane. For those that lie in the convex hull, we calculate the signed distance from the support surface (the positive is in the direction of \mathbf{n}_p), keeping finally only those with positive distance. For these points, we apply a Euclidean clustering to extract the object clouds on the table. In our scenario, we keep the largest cluster as the object to be grasped. Note that in this stage, any object detection method can be applied, if a particular object needs to be grasped; also other types of clustering (e.g., using normals and curvatures) can be used in the place of the Euclidean to improve the segmentation.

For the first stage of the algorithm, to find the visually estimated CoM position, a 3D voxel grid of the object point cloud is created. Each voxel is of fixed size and for each one, we replace the set of points that lies in it with their centroid. The 3D voxelization is needed to be able to distribute equally the acquired points on the object. Then, the CoM position is simply the median of the voxelized object point cloud. Note that the estimated CoM may either be on a graspable area, a non-graspable one, or even outside the object. For this reason, we also need to localize all the possible grasps on the object and select the closest to the CoM one. This part of the method geometrically computes a CoM position, which is necessary for the initial object grasp.

A set of uniformly distributed cylindrical (quadratic curve) grasps are localized in real-time on the object of the size of the hand as described in Sec. 2. The amount of grasps oversamples the object in a way such that there exists at least one grasp per centimetre in the graspable areas of the object. A small cylindrical gap without points is guaranteed from the method, to accommodate for the grasping. This part of the method makes it generic to all the objects that include handles. From all these handle-like grasps, the closest to the estimated CoM is selected. Its frame $C_f = (\hat{\mathbf{x}}_c, \hat{\mathbf{y}}_c, \hat{\mathbf{z}}_c)$ is defined as follows. The $\hat{\mathbf{x}}_c$ -axis is the cylinder axis pointing to the right, the $\hat{\mathbf{y}}_c$ -axis is the unit normal vector \mathbf{n}_p , while the $\hat{\mathbf{z}}_c$ -axis is uniquely defined as the cross-product between the $\hat{\mathbf{x}}_c$ -axis and the $\hat{\mathbf{y}}_c$ -axis towards the range sensor. The origin point \mathbf{c} , which is initially defined as the center of the cylinder, is translated on the surface of the local point cloud neighborhood in the direction of the normal vector \mathbf{n}_p . This CoM grasp frame C_f at its origin \mathbf{c} is the one that needs to co-align with the hand frame H_f (see Fig. 3(b)) during the grasping.

Point cloud filtering is an important step to make the fitting method⁶⁴ work on our stereo camera system. The method was originally developed for the very accurate structured light Asus XTion sensor, which preserves the cylindrical geometry of a surface. In contrary, our stereo camera point cloud is showing a big number of outliers and local spikes. For this reason, a real-time statistical outlier removal and

a second degree moving least squares filtering has been applied on the object point cloud.

3.2. Proprioceptive-based grasp estimation

A visual CoM localization from an exteroceptive range sensor could provide a first grasp estimation. Due to the point cloud data uncertainties (e.g., variations in points position, outliers, or missing areas due to occlusions) and a potential uneven distribution of mass along the object, the visually estimated CoM position may not be the same as the actual one. Moreover, an exteroceptive range sensor is limited to represent only points on the surface of an object. Using the F/T sensor, which is installed at the wrist level, a 6 DoF force and moment vector can be measured. From these and the vision in the loop, the 3D displacement can be calculated, through a sequence of grasps and lifts, such that the wrist torque is minimized. In particular, we follow the next steps, called ‘Stage II’ (illustrated also as a flowchart, presented in Fig. 5):

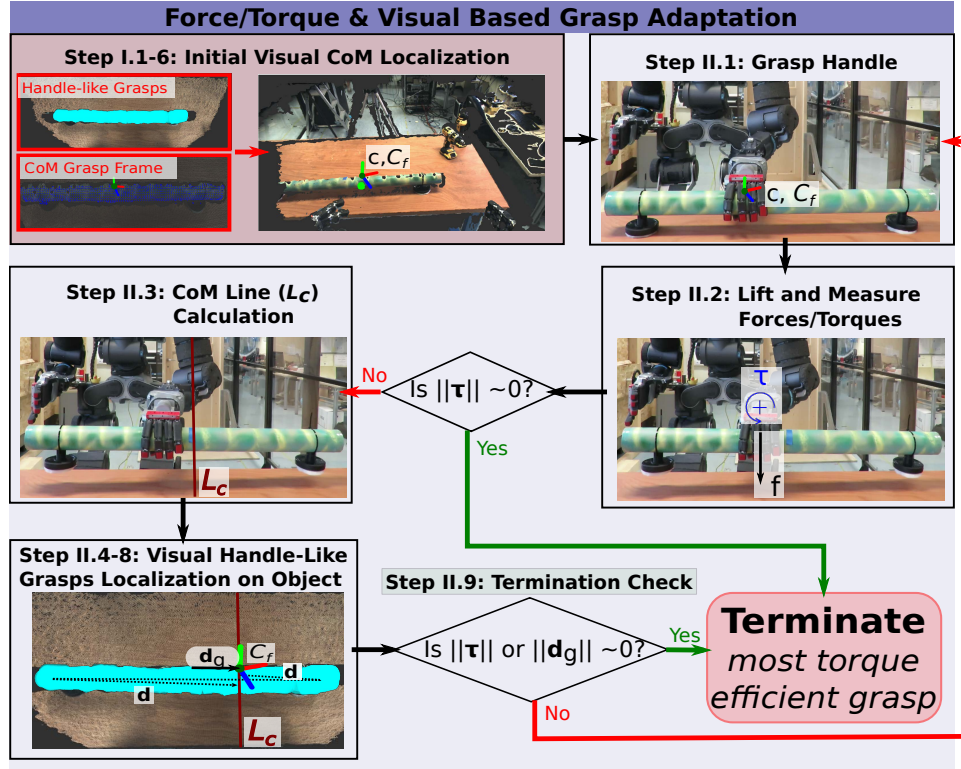


Fig. 5: The flowchart of the whole algorithm: the initial 3D range-based localization of the CoM (Stage I) and the wrist force/torque-based regrasping adaptation (Stage II).

Stage II (force/torque-based grasp adaptation)

- **Step II.1 [grasp handle]:** Approach and grasp the object at the selected grasp handle.
- **Step II.2 [lift and measure forces/torques]:** Lift the object slightly and measure the forces and torques from the wrist F/T sensor. If the minimum torque threshold has been reached, terminate. Otherwise, lower and release the object.
- **Step II.3 [CoM line (L_c) calculation]:** Based on the forces and torques, calculate the *CoM line* that goes through the CoM point of the object.
- **Step II.4–8 [visual handle-like grasps localization on object]** Run Steps I.1–4 (Sec. 3.1). Select as the next handle grasp the one with the minimum torsional effort with respect to the CoM line.
- **Step II.9 [termination check]** If the minimum displacement distance or minimum torque has been reached, terminate after grasping and lifting the object. Run Step I.6 (Sec. 3.1) and go to Step II.1.

3.2.1. The algorithm

A final grasp may bring grasping instabilities for heavy objects with irregular mass distribution, when the actual CoM is far from the grasp point estimated only using the object geometry. This is especially true when underactuated hands are used, like those of our robot, where object slips are unavoidable. For this reason, the use of proprioceptive sensing may be essential to improve the initial estimation after the first visually driven grasp.

First, note that from the robot's kinematics, all the data can be transformed to a fixed world frame, and in the rest of the paper, it will be considered to be the Waist frame of the robot. The initial input of this stage is the CoM point and the corresponding closest grasp frame $C_f = (\hat{\mathbf{x}}_c, \hat{\mathbf{y}}_c, \hat{\mathbf{z}}_c)$ at the origin position \mathbf{c} , that were estimated using the exteroceptive 3D visual perception method (Stage I in Sec. 3.1). Initially, the hand approaches and grasps the visual contact point by co-aligning the hand frame H_f with the grasp frame C_f at \mathbf{c} (Fig. 1(d)). Then, the object is lifted slightly from the support surface.

While the object is lifted and it is not moving, the force $\mathbf{f} = (f_x, f_y, f_z)$ and the torque $\boldsymbol{\tau} = (\tau_x, \tau_y, \tau_z)$ vectors are measured at the wrist F/T sensor (their values are averaged over time for two seconds). Based on the standard force/torque relation ($\boldsymbol{\tau} = \mathbf{d} \times \mathbf{f}$), using the property of vector triple product, the distance vector \mathbf{d} to the object's CoM is:

$$\mathbf{d} = \lambda \mathbf{f} + \frac{\mathbf{f} \times \boldsymbol{\tau}}{\|\mathbf{f}\|^2}, \forall \lambda \in \mathbb{R} \quad (1)$$

where $\|\cdot\|$ is the vector norm. This solution represents a set of vectors that go through the CoM point of the object and form a line L_c parallel to the gravity

vector. If the torque effort, after the object is lifted, is smaller than a threshold, i.e., $\|\boldsymbol{\tau}\| \leq \tau_{thres}$, it is assumed that the real CoM has been reached and no further action is required.

Otherwise, given the CoM line L_c , the object is lowered and released on the support surface and the visual stage is repeated as follows. The object is segmented from the support surface and a new set of handle-like grasps are localized on it (Steps I.1-4). Then, these handle grasps are evaluated with respect to the line L_c and their potential torque effort, such that the most wrist torque efficient one is selected. To do so, the displacement vector \mathbf{d}_h between each handle-like grasp frame (computed in Step I/II.6) and the CoM line L_c needs to be calculated. The problem of calculating the distance between a point and a line is a standard calculus problem.⁵ Assuming that the applied force \mathbf{f} is only due to the object's mass towards the gravity, which was computed during the first object lift, the handle frame with the minimum potential torque effort, defined as the norm of the torque vector (i.e., $\|\boldsymbol{\tau}_h\| = \|\mathbf{d}_h \times \mathbf{f}\|$), is selected. Moreover, the corresponding displacement vector \mathbf{d}_g to the new CoM estimation and its length $\|\mathbf{d}_g\|$ are extracted.

If the new grasp frame displacement is smaller than a threshold distance d_{thres} , i.e., $\|\mathbf{d}_g\| \leq d_{thres}$, it is assumed that the closest grasp to the real CoM point has been achieved and no further action is required. Otherwise, the object is lowered and regasped in the new displaced position, starting the second stage loop from the beginning. The displacement threshold is required given that an object may not have a feasible grasp close to its CoM, and thus the torque effort threshold will not be efficient for the termination.

4. EXPERIMENTS

To test the overall approach, we run experiments on the humanoid robot WALK-MAN (introduced in Sec. 2.2). We first test the ability to visually detect the CoM on various objects on a table. Then, we test the regrasping process on three types of objects, by also changing their mass distribution: (i) a handle object that the real CoM is along its handle (Fig. 6—first column), (ii) a more complex object that the real CoM is outside the object and includes non-graspable areas (Fig. 6—second column), and finally (iii) an object that its CoM is inside the object, but in a non-graspable position (Fig. 6—third column). We next discuss the hardware, control, and planning setup, as well as the experimental apparatus with the results.

4.1. Control and Planning System

The robot is controlled with the XBotCore⁴⁷ and the YARP middleware framework,⁴⁴ while all the visual and force/torque perception data are handled with ROS. Using the YARP functionalities, the high-level commands are created for the required motion primitives (e.g., “reach”, “grasp”, “lift”, etc.) and delivered to the low-level torque controller, implemented on DSPs at each joints.⁶⁰ In particular, to control the whole-body motion of the robot, inverse kinematics is resolved by

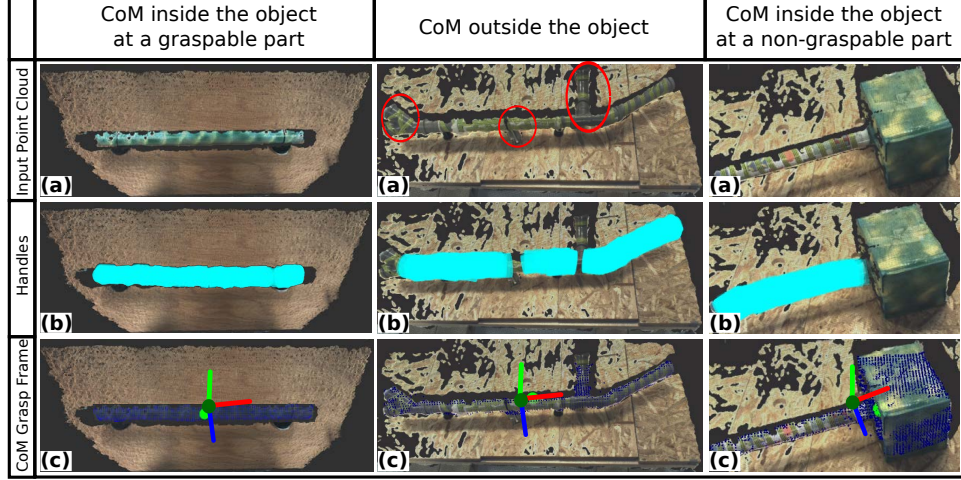


Fig. 6: The method for three object types: a handled one with its CoM at a graspable point, a partially handled one with its CoM outside the object, and one with its CoM at a non-graspable position.

the Stack-of-Task (SoT) formalism,¹³ which employs cascaded Quadratic Programming (QP) solvers to efficiently find an optimum, in the least-square sense with a description of hierarchical tasks and constraints. The OpenSoT control library has been used to provide these features.⁴⁵ Throughout the experiment, a single arm is controlled to handle the object manipulation depending on the grasp position. The position of the other arm and the lower body are regulated, while the CoM of the full body is controlled to reside in the convex hull, for stable balancing during the task.

4.2. Experimental Apparatus

For the exteroceptive experimental testing, we ran the 3D range CoM estimation on various objects (handled or not) and evaluated qualitatively the results, some of which appear in Fig. 7.

For the CoM-based grasp adaptation experiments, we set the robot in home position, 50cm in front of a flat 90cm-tall table, where we place the objects in a reachable distance (see Fig. 1—upper left). As shown in Fig. 8, the first object (Exps. Nos. 1–5^a; Fig. 6—first column) is a cylindrical debris of 7cm diameter and 105cm long, on which we attach 0.5kg weights. First, we attach two 0.5kg weights, 10cm distanced from each end (Exp. No. 1) and for each experiment we move the left weight 10cm right (Exps. Nos. 1–4), changing in this way the position of the real CoM on the object (red dot). We then add a third 0.5kg weight (Exp. No. 5) and after the first regrasping we remove it manually to test the real-time online CoM position reestimation and the success of the regrasping according to the new

^aOnly one experiment is visualized, while the rest can be found in the videos.

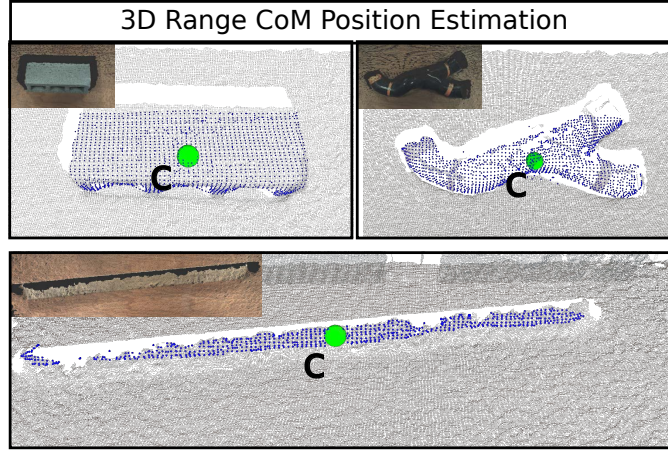


Fig. 7: A set of three objects and the estimated CoM (green dot) using 3D range voxelization.

measurements. The second object (Exps. Nos. 6–9^a; Fig. 6—second column) is a set of connected 5cm diameter cylindrical parts. Only some of them are graspable, while its actual CoM is not inside the object. For each experiment we attach a 0.25kg weight that each time we move it 10cm left, along the white handle of the object. The third object (Exp. No. 10; Fig. 6—third column) is a hammer-like one, where the CoM is inside its rectangle. Last but not the least, we also tried our method on a small-handled object, i.e., a drill (Exp. No. 11).

Evaluation Measures

For the first five experiments, we recorded the real CoM position d_r , the visually calculated one d_v (measured from the left most part of the object), the force f_g along the gravity vector, the torque norm $\tau = \|\tau\|$ after the initial 3D visual grasp, the computed displacement distance $d = \|\mathbf{d}_g\|$, the new torque τ' , and the displacement value d' after the new pose regrasp. For the rest of the experiments (Exps. Nos. 6–10) we recorded the initial torque τ and the final one τ' after the regrasp. Each experiment is performed 10 times and Tables 1 and 2 present the average recorded measurements, while Figs. 9 and 10 visualize the CoM position deviations and torques for the initial proprioceptive and the reestimated exteroceptive grasplings. Note that in the beginning of each experiment, we remove all the force/torque sensor residuals, before grasping the object, while we manually set the thresholds $d_{thres} = 2\text{cm}$ and $\tau_{thres} = 0.02\text{Nm}$. All the experimental videos can be found under the following link:

http://dkanou.github.io/projects/com_grasping

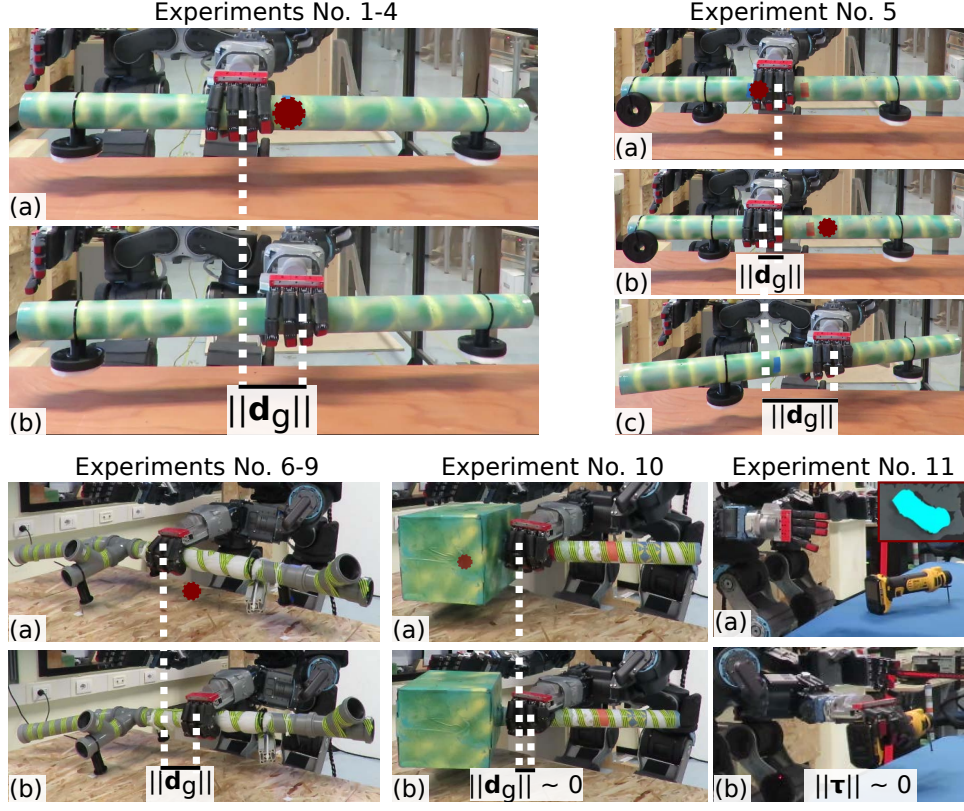


Fig. 8: Experimental instances of: (1) a handled object and the moving weights along its axis (Exps. Nos. 1–4); (2) a handled object and the removed weight from its left part (Exp. No. 5); (3) a partially-handled object, whose CoM is outside the object (Exps. Nos. 6–9); (4) an object, whose CoM is at a non-graspable area (Exp. No. 10). Every sub-figure (a) visualizes the initial visual-based grasp and sub-figure (b) the minimum torque regrasp adaptation, while for Exp. No. 5, sub-figure (c) shows the regrasp adaptation after a weight is removed. Exp. No. 11 shows an experimental instance of a small handled object (drill), with the (a) localized handles and (b) its minimum torque initial visual-based grasp. The red star denotes the real CoM position.

Table 1: Average results of the first five experiments (Exps. Nos. 1–5).

No.	real d_r (m)	3D visual CoM grasp				F/T CoM regrasp		
		d_v (m)	f_g (N)	τ (Nm)	d (m)	τ' (Nm)	d' (m)	$ \Delta\tau/\tau $ (%)
1	0.520	0.450	12.0	0.886	0.0736	-0.031	-0.0020	96.5
2	0.565	0.480	13.7	1.323	0.0969	0.090	0.0069	93.2
3	0.600	0.490	15.9	1.734	0.1086	0.178	0.0177	89.7
4	0.645	0.500	14.2	1.617	0.1138	0.147	0.0112	90.9
5	0.360	0.390	17.9	-0.648	-0.0362	0.148	0.0084	22.8
	0.530	(0.390)	14.9	1.514	0.1013	0.112	0.0082	92.6

Table 2: Average results of the last five experiments (Exps. Nos. 6–10).

No.	Visual grasp	F/T regrasp	
	τ (Nm)	τ' (Nm)	$ \Delta\tau/\tau $ (%)
6	0.460	0.1724	62.5
7	1.162	0.7075	39.1
8	1.096	0.7561	31.0
9	1.079	0.7153	33.7
10	1.330	0.538	59.5




Fig. 9: For each of the 10 repetitions of the Exps. Nos. 1–5, the initial visual CoM grasp position d_u , the regrasped adaptation displacement distance ($d_u - d$), and the real CoM position (d_r) are visualized. All distances are measured from the left most part of the object.

4.3. Results

We first note that the visual system, with the handle-like grasp detection and the selection of the closest to the CoM estimation one, is working very reliably. There was never noted any failure in the grasping.

With the first five experiments (Exps. Nos. 1–5) we tested the ability of the method to reach the real CoM when this is inside the object. The particular cylindrical object is everywhere graspable. From the average results of the first four experiments (Exps. Nos. 1–4) in Table 1 and Figs. 9 and 10, we first note that the visually computed CoM is always found roughly in the center of the object, which gives a reasonable initial grasp point. One regrasp using the force/torque perception method was always enough to reach the d_{thres} threshold from the CoM or the τ_{thres} threshold when it is lifted, making the final torque (τ') or the displacement (d') very close to zero. The percentage of torque improvement (last column of Table 1 and

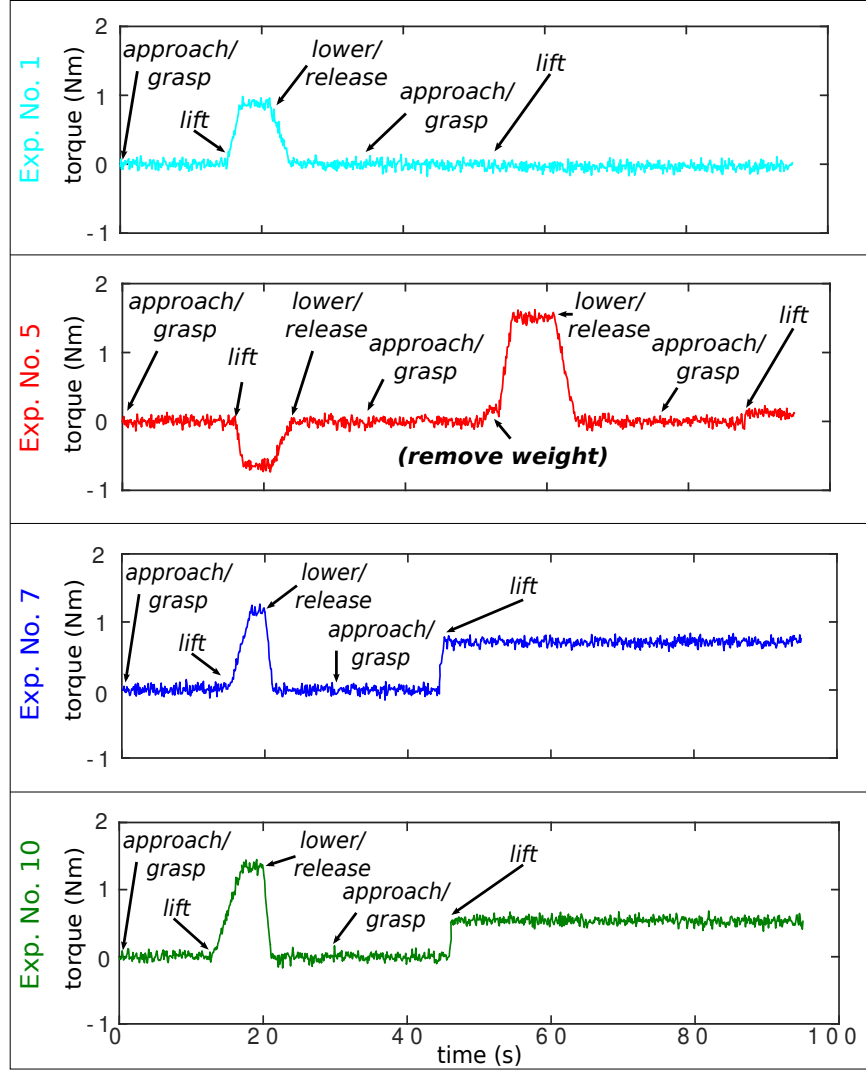


Fig. 10: The continuous wrist torque measurements during the Exps. No. 1, 5, 7, and 10 are presented. The values correspond to Table 1 and Table 2 ones. The torque values for the rest of the experiments are omitted since they are similar to the presented ones. In each sub-graph, the actions of the robot on the object are labelled as: "approach/grasp", "lift", and "lower/release", to show the torque value changes during each experiment.

Fig. 10) is mostly high. Together with the low final torques it means that the second grasp is always more wrist torque efficient and reliable, since the object is grasped very close to the real CoM. In the fifth experiment (Exp. No. 5) we verified that our method can detect online changes in the distribution of the object mass and automatically regrasp the object at its newly estimated CoM position using the force and torque readings.

In the next four experiments (Exps. Nos. 6–9; Table 2 and Fig. 10) we tested

the ability of the method to reach the closest possible grasp to the real CoM (where the torque is the minimum possible using a single hand), when the CoM is out of the object and the object includes non-graspable areas. We note that the torque is always minimized after the regrasp, which makes the grasping more torque efficient. In the tenth experiment we tested the ability of finding the closest grasp to the CoM that is inside the object but in a non-graspable position. The final grasp is close to the first one. Note that the slight hand displacement (very close to zero) is due to different set of handle-like grasps localization during the visual stage of the second regrasp iteration.

Last but not least, in the eleventh experiment (Exp. No. 11) we just tested the ability to recognize a negligible torque during the first object lift, since the object (drill) is very small. This is the reason that a regrasp is not required. In all cases the possible minimum torque was reached after a single regrasp.

4.4. Discussion and Limitations

The results of the proposed framework show that the combination of proprioceptive 3D range and exteroceptive force/torque sensing decreases the torque effort at the wrist level, when lifting heavy objects with one hand. The methodology, as appears in this paper, has some limitation that we discuss briefly in this section.

To start with the proprioceptive 3D range part, the presented method segments an object from the environment, assuming a dominant plane (e.g., a table) that the objects lies on. The extracted 3D point cloud of the object is then used to estimate visually the object's CoM. This simplified segmentation method has an obvious limitation in cases where the objects do not lie on a table, e.g., objects in a bin, or when multiple ones are occluded. Given though that this part acts as a black box in our framework, one can apply any object segmentation method (including object recognition/detection and localization) to extract the required point cloud in real-world environments. In addition, during voxelization for localizing the CoM, we assumed that the object is made by isotropic material with constant mass density. Given that for some objects, such as hammers, this assumption may not hold, a more sophisticated technique should be developed to detect different material densities of objects and adapt accordingly the voxelization technique. Furthermore, the focus of this paper is on objects than include handle-like graspable areas. As previously mentioned in the related work (Sec. 1.1), newly developed methods can localize various types of graspable areas on robots (using for instance deep learning). We envision a system that includes such methods, but given that a big set of objects have handles, in this paper, we focused primarily on them.

During the exteroceptive force/torque regrasping part, the method minimizes the torque effort at the wrist level of the robot, where the sensor is installed. In this way, the produced torque in the rest of the robot's joints is not considered. Moreover, the robot does not change the orientation of the object while lifting it. These types of maneuvers could minimize the torque effort at the wrist level. In

addition, the method considers only a single arm use, while the framework has more potentials using two arms or even the whole body to minimize the torque effort. Last but not the least, “full palm” grasps limit the introduced framework in the type of hands that can be used. Other types, such as precision grippers, can introduce slippage during object lifting, resulting in potential torque imprecisions. Thus, the method needs to be extended to monitor torques and slippage during lifting, for instance, with the use of tactile sensing.

Experimentally, the method has been tested on relatively simplified objects, whereas the experiments should be extended to more complex ones (such as chairs), where occlusions and grasping limitations may exist.

5. Conclusions and Future Works

In this paper, we presented a novel combination of 3D range and force/torque sensing for finding CoM-based grasps on heavy objects that include handles. By first applying a visual CoM estimation using point cloud data coming from a range sensor and then applying a set of regrasps to measure the forces and torques on the wrist of the arm, our method is able to accurately detect the real CoM position of the object and grasp it from the most torque efficient handle grasp. In the experiments, we showed that one regrasp after the visual-based one is enough to localize the most wrist torque efficient one on the humanoid robot WALK-MAN.

In future work, we first plan to improve the visual estimation of the CoM by using a SLAM method like the Moving KinectFusion⁵⁸ for building a better point cloud representation of the object while the head or the whole robot is moving. We also plan to generalize our method by using two hands for the manipulation or by considering whole-body motions for more secure bi-manual grasping. One further extension could be the application of different strategies when torque/force limits are reached during the object lifting phase, since the object may slip or rotate during a hand grasp.

Acknowledgment

This work is supported by the FP7-ICT-2013-10 WALK-MAN European Commission project, no. 611832.

References

1. Anna A. Petrovskaya and Oussama Khatib, *Global Localization of Objects via Touch*, IEEE Transactions on Robotics (T-RO) **27** (2011), no. 3, 569–585.
2. Hussam Al Hussein, Tiago Caldeira, Dongming Gan, Jorge Dias, and Lakmal Seneviratne, *Object Shape Perception in Blind Robot Grasping using a Wrist Force/Torque Sensor*, IEEE 20th International Conference on Electronics, Circuits, and Systems (ICECS) (Abu Dhabi, United Arab Emirates), IEEE, 2013, pp. 193–196.
3. Peter K. Allen, Andrew T. Miller, Paul Y. Oh, and Brian S. Leibowitz, *Integration of Vision, Force and Tactile Sensing for Grasping*, International Journal of Intelligent Machines **4** (1999), no. 1, 129–149.

4. Christopher G. Atkeson, Chae H. An, and John M. Hollerbach, *Rigid Body Load Identification for Manipulators*, IEEE 24th Conference on Decision and Control (Fort Lauderdale, FL, USA), IEEE, 1985, pp. 996–1002.
5. John Perry Ballantine and Arthur Rudolph Jerbert, *Distance from a Line or Plane to a Point*, American Mathematical Monthly **59** (1952), no. 4, 242–243.
6. Yasemin Bekiroglu, Renaud Detry, and Danica Kragic, *Learning Tactile Characterizations of Object- and Pose-Specific Grasps*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (San Francisco, CA, USA), IEEE/RSJ, 2011, pp. 1554–1560.
7. Yasemin Bekiroglu, Dan Song, Lu Wang, and Danica Kragic, *A Probabilistic Framework for Task-Oriented Grasp Stability Assessment*, IEEE International Conference on Robotics and Automation (ICRA) (Karlsruhe, Germany), IEEE, 2013, pp. 3040–3047.
8. Antonio Bicchi and Vijay Kumar, *Robotic Grasping and Contact: a Review*, IEEE International Conference on Robotics and Automation (ICRA) (San Francisco, CA, USA), IEEE, 2000, pp. 348–353.
9. Joao Bimbo, Lakmal D. Seneviratne, Kaspar Althoefer, and Hongbin Liu, *Combining Touch and Vision for the Estimation of an Object's Pose During Manipulation*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Tokyo, Japan), IEEE/RSJ, 2013, pp. 4021–4026.
10. Manuel G. Catalano, Giorgio Grioli, Alessandro Serio, Edoardo Farnioli, Cristina Piazza, and Antonio Bicchi, *Adaptive Synergies for a Humanoid Robot Hand*, IEEE-RAS 12th International Conference on Humanoid Robots (Humanoids) (Osaka, Japan), IEEE-RAS, 2012, pp. 7–14.
11. Lillian Chang, Joshua R. Smith, and Dieter Fox, *Interactive Singulation of Objects from a Pile*, IEEE International Conference on Robotics and Automation (ICRA) (Saint Paul, MN, USA), IEEE, 2012, pp. 3875–3882.
12. Renaud Detry, Carl Henrik Ek, Marianna Madry, and Danica Kragic, *Learning a Dictionary of Prototypical Grasp-Predicting Parts from Grasping Experience*, IEEE International Conference on Robotics and Automation (ICRA) (Karlsruhe, Germany), IEEE, 2013, pp. 601–608.
13. Adrien Escande, Nicolas Mansard, and Pierre-Brice Wieber, *Hierarchical Quadratic Programming: Fast Online Humanoid-Robot Motion Generation*, International Journal of Robotics Research (IJRR) **33** (2014), no. 7, 1006–1028.
14. David Fischinger and Markus Vincze, *Empty the Basket - a Shape Based Learning Approach for Grasping Piles of Unknown Objects*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Vilamoura, Portugal), IEEE/RSJ, 2012, pp. 2051–2057.
15. David Fischinger, Markus Vincze, and Yun Jiang, *Learning Grasps for Unknown Objects in Cluttered Scenes*, IEEE International Conference on Robotics and Automation (ICRA) (Karlsruhe, Germany), IEEE, 2013, pp. 609–616.
16. Marcus Gualtieri, Andreas ten Pas, Kate Saenko, and Robert Platt, *High Precision Grasp Pose Detection in Dense Clutter*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Daejeon, South Korea), IEEE/RSJ, 2016, pp. 598–605.
17. Paul Hebert, Nicolas Hudson, Jeremy Ma, and Joel Burdick, *Fusion of Stereo Vision, Force-Torque, and Joint Sensors for Estimation of In-Hand Object Location*, IEEE International Conference on Robotics and Automation (ICRA) (Shanghai, China), IEEE, 2011, pp. 5935–5941.
18. Alexander Herzog, Peter Pastor, Mrinal Kalakrishnan, Ludovic Righetti, Tamim Asfour, and Stefan Schaal, *Template-Based Learning of Grasp Selection*, IEEE Inter-

- national Conference on Robotics and Automation (ICRA) (Saint Paul, MN, USA), IEEE, 2012, pp. 2379–2384.
19. Dirk Holz, Stefan Holzer, Radu Bogdan Rusu, and Sven Behnke, *Real-Time Plane Segmentation using RGB-D Cameras*, 15th RoboCup International Symposium (Istanbul, Turkey), Lecture Notes in Computer Science, vol. 7416, Springer, July 2011, pp. 307–317.
 20. Stefan Holzer, Radu Bogdan Rusu, Michael Dixon, Suat Gedikli, and Nassir Navab, *Adaptive Neighborhood Selection for Real-Time Surface Normal Estimation from Organized Point Cloud Data Using Integral Images*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Vilamoura, Portugal), IEEE/RSJ, 2012, pp. 2684–2689.
 21. Kyuhei Honda, Tsutomu Hasegawa, Toshihiro Kiriki, and Takeshi Matsuoka, *Real-time Pose Estimation of an Object Manipulated by Multi-Fingered Hand using 3D Stereo Vision and Tactile Sensing*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Victoria, BC, Canada), IEEE/RSJ, 1998, pp. 1814–1819.
 22. Lorenzo Jamone, Lorenzo Natale, Giorgio Metta, and Giulio Sandini, *Highly Sensitive Soft Tactile Sensors for an Anthropomorphic Robotic Hand*, IEEE Sensors Journal **15** (2015), no. 8, 4226–4233.
 23. Yun Jiang, Stephen Moseson, and Ashutosh Saxena, *Efficient Grasping from RGBD Images: Learning using a New Rectangle Representation*, IEEE International Conference on Robotics and Automation (ICRA) (Shanghai, China), IEEE, 2011, pp. 3304–3311.
 24. Peter Kaiser, Eren E. Aksoy, Markus Grotz, Dimitrios Kanoulas, Nikos G. Tsagarakis, and Tamim Asfour, *Experimental Evaluation of a Perceptual Pipeline for Hierarchical Affordance Extraction*, 2016 International Symposium on Experimental Robotics (ISER), vol. 1, Springer, 2017, pp. 136–146.
 25. Peter Kaiser, Dimitrios Kanoulas, Markus Grotz, Luca Muratore, Alessio Rocchi, Enrico Mingo Hoffman, Nikos G. Tsagarakis, and Tamim Asfour, *An Affordance-Based Pilot Interface for High-Level Control of Humanoid Robots in Supervised Autonomy*, IEEE-RAS International Conference on Humanoid Robots (Humanoids) (Cancun, Mexico), IEEE-RAS, 2016, pp. 621–628.
 26. Dimitrios Kanoulas, *Curved Surface Patches for Rough Terrain Perception*, Ph.D. thesis, CCIS, Northeastern University, August 2014.
 27. Dimitrios Kanoulas, Jinoh Lee, Darwin G. Caldwell, and Nikos G. Tsagarakis, *Visual Grasp Affordance Localization in Point Clouds using Curved Contact Patches*, International Journal of Humanoid Robotics (IJHR) **14** (2017), no. 1, 1650028–1–1650028–21.
 28. Dimitrios Kanoulas and Marsette Vona, *The Surface Patch Library (SPL)*, IEEE International Conference on Robotics and Automation (ICRA) Workshop: MATLAB/Simulink for Robotics Education and Research (Hong Kong), IEEE, 2014, dikanou.github.io/projects/spl/, pp. 1–9.
 29. Dov Katz, Moslem Kazemi, J. Andrew Bagnell, and Anthony Stentz, *Clearing a Pile of Unknown Objects using Interactive Perception*, 2013 IEEE International Conference on Robotics and Automation (ICRA) (Karlsruhe, Germany), IEEE, 2013, pp. 154–161.
 30. Arie Kaufman, Daniel Cohen, and Roni Yagel, *Volume Graphics*, Computer **26** (1993), no. 7, 51–64.
 31. Moslem Kazemi, Jean-Sebastien Valois, J. Andrew Bagnell, and Nancy Pollard, *Robust Object Grasping using Force Compliant Motion Primitives*, Robotics: Science and Systems (RSS) (Sydney, Australia), MIT Press, 2012, pp. 177–185.

32. Charles C. Kemp, Aaron Edsinger, and Eduardo Torres-Jara, *Challenges for Robot Manipulation in Human Environments [Grand Challenges of Robotics]*, IEEE Robotics and Automation Society RAM **14** (2007), no. 1, 20–29.
33. Ellen Klingbeil, Deepak Rao, Blake Carpenter, Varun Ganapathi, Andrew Y. Ng, and Oussama Khatib, *Grasping with Application to an Autonomous Checkout Robot*, IEEE International Conference on Robotics and Automation (ICRA) (Shanghai, China), IEEE, 2011, pp. 2837–2844.
34. Oliver Kroemer and Jan Peters, *Predicting Object Interactions from Contact Distributions*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Chicago, IL, USA), IEEE/RSJ, 2014, pp. 3361–3367.
35. Joanna Laaksonen, Ekaterina Nikandrova, and Ville Kyrki, *Probabilistic Sensor-Based Grasping*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Vilamoura, Portugal), IEEE/RSJ, 2012, pp. 2019–2026.
36. Steven M. LaValle, *Planning algorithms*, Cambridge University Press, New York, NY, USA, 2006.
37. Ian Lenz, Honglak Lee, and Ashutosh Saxena, *Deep Learning for Detecting Robotic Grasps*, The International Journal of Robotics Research (IJRR), Special Issue on Robot Vision **34** (2015), no. 4-5, 705–724.
38. Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen, *Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection*, CoRR **abs/1603.02199** (2016).
39. Miao Li, Yasemin Bekiroglu, Danica Kragic, and Aude Billard, *Learning of Grasp Adaptation through Experience and Tactile Sensing*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Chicago, IL, USA), IEEE/RSJ, 2014, pp. 3339–3346.
40. Rui Li, Robert Platt Jr., Wenzhen Yuan, Andreas ten Pas, Nathan Roscup, Mandayam A. Srinivasan, and Edward Adelson, *Localization and Manipulation of Small Parts using GelSight Tactile Sensing*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Chicago, IL, USA), IEEE/RSJ, 2014, pp. 3988–3993.
41. Efraim Lopez-Damian, Daniel Sidobre, and Rachid Alami, *A Grasp Planner Based On Inertial Properties*, IEEE International Conference on Robotics and Automation (ICRA) (Barcelona, Spain), IEEE, 2005, pp. 754–759.
42. Tanis Mar, Vadim Tikhonoff, Giorgio Metta, and Lorenzo Natale, *Self-supervised Learning of Grasp Dependent Tool Affordances on the iCub Humanoid Robot*, IEEE International Conference on Robotics and Automation (ICRA) (Seattle, WA, USA), IEEE, 2015, pp. 3200–3206.
43. Matthew T. Mason and J. Kenneth Salisbury Jr., *Robot Hands and the Mechanics of Manipulation*, MIT Press, 1985.
44. Giorgio Metta, Paul Fitzpatrick, and Lorenzo Natale, *YARP: Yet Another Robot Platform*, International Journal on Advanced Robotics Systems **3** (2006), no. 1, 43–48.
45. Enrico Mingo Hoffman, Alessio Rocchi, Arturo Laurenzi, and Nikos G. Tsagarakis, *Robot control for dummies: Insights and examples using OpenSoT*, IEEE-RAS International Conference on Humanoid Robots (Humanoids) (Birmingham, UK), IEEE-RAS, 2017, pp. 736–741.
46. Luis Montesano, Manuel Lopes, Alexandre Bernardino, and Jose Santos-Victor, *Learning Object Affordances: From Sensory-Motor Coordination to Imitation*, IEEE Transactions on Robotics (T-RO) **24** (2008), no. 1, 15–26.
47. Luca Muratore, Arturo Laurenzi, Enrico Mingo Hoffman, Alessio Rocchi, Darwin G. Caldwell, and Nikos G. Tsagarakis, *XBotCore: A Real-Time Cross-Robot Software*

- Platform*, IEEE International Conference on Robotic Computing (IRC) (Taichung, Taiwan), IEEE, 2017, pp. 77–80.
48. Richard M. Murray, S. Shankar Sastry, and Li Zexiang, *A Mathematical Introduction to Robotic Manipulation*, CRC Press, Inc. Boca Raton, FL, USA, 1994.
 49. Austin Myers, Ching L. Teo, Cornelia Fermuller, and Yiannis Aloimonos, *Affordance Detection of Tool Parts from Geometric Features*, IEEE International Conference on Robotics and Automation (ICRA) (Seattle, WA, USA), IEEE, 2015, pp. 3200–3206.
 50. Anh Nguyen, Dimitrios Kanoulas, Darwin G. Caldwell, and Nikos G. Tsagarakis, *Detecting Object Affordances with Convolutional Neural Networks*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Daejeon, South Korea), IEEE/RSJ, 2016, pp. 2765–2770.
 51. ———, *Preparatory Object Reorientation for Task-Oriented Grasping*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Daejeon, South Korea), IEEE/RSJ, 2016, pp. 893–899.
 52. ———, *Object-Based Affordances Detection with Convolutional Neural Networks and Dense Conditional Random Fields*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Vancouver, BC, Canada), IEEE/RSJ, 2017, pp. 5908–5915.
 53. Allison M. Okamura and Mark R. Cutkosky, *Feature Detection for Haptic Exploration with Robotic Fingers*, The International Journal of Robotics Research (IJRR) **20** (2001), no. 12, 925–938.
 54. Anna Petrovskaya and Kaijen Hsiao, *Active Manipulation for Perception*, pp. 1037–1062, Springer International Publishing, Cham, 2016.
 55. Anna Petrovskaya, Oussama Khatib, Sebastian Thrun, and Andrew Y. Ng, *Bayesian Estimation for Autonomous Object Manipulation Based on Tactile Sensors*, IEEE International Conference on Robotics and Automation (ICRA) (Orlando, FL, USA), IEEE, 2006, pp. 707–714.
 56. Alessandro Pieropan, Carl Henrik Ek, and Hedvig Kjellström, *Functional Object Descriptors for Human Activity Modeling*, IEEE International Conference on Robotics and Automation (ICRA) (Karlsruhe, Germany), IEEE, 2013, pp. 1282–1289.
 57. Lerrel Pinto and Abhinav Gupta, *Supersizing Self-supervision: Learning to Grasp from 50K Tries and 700 Robot Hours*, IEEE International Conference on Robotics and Automation (ICRA) (Stockholm, Sweden), IEEE, 2016, pp. 3406–3413.
 58. Henry Roth and Marsette Vona, *Moving Volume KinectFusion*, British Machine Vision Conference (BMVC) (Surrey, UK), BMVA Press, Sept. 2012, pp. 1–11.
 59. Radu Bogdan Rusu and Steve Cousins, *3D is here: Point Cloud Library (PCL)*, IEEE International Conference on Robotics and Automation (ICRA) (Shanghai, China), IEEE, 2011, pp. 1–4.
 60. Alessandro Settimi, Corrado Pavan, Valerio Varricchio, Mirko Ferrati, Enrico Mingo Hoffman, Alessio Rocchi, Kamilo Melo, Nikos G. Tsagarakis, and Antonio Bicchi, *A Modular Approach for Remote Operation of Humanoid Robots in Search and Rescue Scenarios*, International Workshop on Modelling and Simulation for Autonomous Systems (MESAS), vol. 8906, Springer, Switzerland, 2014, pp. 192–205.
 61. Michael Stark, Philipp Lies, Michael Zillich, Jeremy Wyatt, and Bernt Schiele, *Functional Object Class Detection Based on Learned Affordance Cues*, 6th International Conference Computer Vision Systems (ICVS) (Berlin, Heidelberg), Springer, 2008, pp. 435–444.
 62. Mike Stilman, Koichi Nishiwaki, and Satoshi Kagami, *Learning Object Models for Whole Body Manipulation*, IEEE-RAS 7th International Conference on Humanoid Robots (Humanoids) (Pittsburgh, PA, USA), IEEE-RAS, 2007, pp. 174–179.

63. Jörg Stücker, Ricarda Steffens, Dirk Holz, and Sven Behnke, *Efficient 3D Object Perception and Grasp Planning for Mobile Manipulation in Domestic Environments*, Robotics and Autonomous Systems (RAS) **61** (2013), no. 10, 1106–1115.
64. Andreas ten Pas and Robert Platt, *Localizing Grasp Affordances in 3-D Points Clouds Using Taubin Quadric Fitting*, International Symposium on Experimental Robotics (ISER) (Marrakech and Essaouira, Morocco), Springer, 2014.
65. ———, *Using Geometry to Detect Grasps in 3D Point Clouds*, The International Symposium on Robotics Research (ISRR) (Sestri Levante, Italy), Springer, 2015.
66. Nikos G. Tsagarakis, D. G. Caldwell, F. Negrello, W. Choi, L. Baccelliere, V.G. Loc, J. Noorden, L. Muratore, A. Margan, A. Cardellino, L. Natale, E. Mingo Hoffman, H. Dallali, N. Kashiri, J. Malzahn, J. Lee, P. Kryczka, D. Kanoulas, M. Garabini, M. Catalano, M. Ferrati, V. Varricchio, L. Pallottino, C. Pavan, A. Bicchi, A. Settini, A. Rocchi, and A. Ajoudani, *WALK-MAN: A High-Performance Humanoid Platform for Realistic Environments*, Journal of Field Robotics (JFR) **34** (2017), no. 7, 1225–1259.
67. Karthik Mahesh Varadarajan and Markus Vincze, *AfRob: The Affordance Network Ontology for Robots*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Vilamoura, Portugal), IEEE/RSJ, 2012, pp. 1343–1350.
68. Giulia Vezzani, Ugo Pattacini, Giorgio Battistelli, Luigi Chisci, and Lorenzo Natale, *Memory Unscented Particle Filter for 6-DOF Tactile Localization*, IEEE Transactions on Robotics (T-RO) **33** (2017), no. 5, 1139–1155.
69. Francisco Vina, Yasemin Bekiroglu, Christian Smith, Yiannis Karayiannidis, and Danica Kragic, *Predicting Slippage and Learning Manipulation Affordances through Gaussian Process Regression*, IEEE-RAS International Conference on Humanoid Robots (Humanoids) (Atlanta, GA, USA), IEEE-RAS, 2013, pp. 462–468.
70. Hanna Yousef, Mehdi Boukallel, and Kaspar Althoefer, *Tactile Sensing for Dexterous In-Hand Manipulation in Robotics — A Review*, Sensors and Actuators A: Physical **167** (2011), no. 2, 171 – 187.
71. Li Zhang, Siwei Lyu, and Jeff Trinkle, *A Dynamic Bayesian Approach to Real Time Estimation and Filtering in Grasp Acquisition*, IEEE International Conference on Robotics and Automation (ICRA) (Karlsruhe, Germany), IEEE, 2013, pp. 85–92.



Dimitrios Kanoulas is a Postdoctoral Researcher at the Humanoids & Human Centered Mechatronics (HHCM) lab at the Istituto Italiano di Tecnologia (IIT), working with Dr. Nikos Tsagarakis and Prof. Darwin Caldwell on the perception and learning part of the COMAN, WALK-MAN, CENTAURO, and CogIMon EU-funded research projects. He was the perception team leader during the DARPA Robotics Challenge Finals in

2015. In August 2014, he completed his Ph.D. at the Geometric and Physical Computing (GPC) Lab in the College of Computer and Information Science at Northeastern University, advised by Prof. Marsette Vona. He started out as a member of the Algorithms and Theory group at Northeastern University, advised by Prof. Rajmohan Rajaraman and Ravi Sundaram. During the summer of 2012, he completed an internship at INRIA in France, advised by Prof. Christian Laugier and Dr. Alexandros Makris. He received his B.S. in Computer Engineering and Informatics Department from University of Patras, Greece in 2008, advised by Prof. Paul

Spirakis and Dr. Haralampos Tsaknakis.



Jinoh Lee received the B.S. degree in Mechanical Engineering from Hanyang University, Seoul, South Korea, in 2003 and the M.Sc. and the Ph.D. degrees in Mechanical Engineering from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea, in 2012. Since 2012, he has joined the Department of Advanced Robotics (ADVR), Istituto Italiano di Tecnologia (IIT), Genoa, Italy, as a Postdoctoral Researcher,

involved in projects such as Safe and Autonomous Physical Human-Aware Robot Interaction (SAPHARI) and Whole-body Adaptive Locomotion and Manipulation (WALK-MAN) funded under the European Community's 7th Framework Programme. He is currently a Research Scientist in ADVR, IIT. His professional is about robotics and control engineering which include whole-body manipulation of high degrees-of-freedom humanoids, robust control of highly nonlinear systems, and compliant robotic system control for safe human-robot interaction. Since 2014, Dr. Lee has participated in Technical Committee member of International Federation of Automatic Control (IFAC), TC4.3 Robotics.



Darwin G. Caldwell received his B.Sc. and Ph.D. degrees in Robotics from the University of Hull, in 1986 and 1990, respectively and an M.Sc. in Management in 1996 from the University of Salford. He is currently a Research Director at the Italian Institute of Technology, Genoa, Italy, and an Honorary Professor at the Universities of Sheffield, Manchester, Kings College London, Bangor, and Tianjin (China). He is the author or co-author of

more than 500 academic papers, and has 15 patents. His research interests include innovative actuators, force augmentation exoskeletons, humanoid (iCub, cCub, CO-MAN and WalkMan) and quadrupled robots (HyQ and HyQ2Max), and medical robotics. In 2015, he was elected a Fellow of the Royal Academy of Engineering.



Nikos G. Tsagarakis received his D.Eng. degree in Electrical and Computer Science Engineering in 1995 from the Polytechnic School of Aristotle University, Greece, an M.Sc. degree in Control Engineering in 1997 and in 2000 a Ph.D. in Robotics from the University of Salford, UK. He is a Tenured Senior Scientist at IIT with overall responsibility for Humanoid design & Human Centred Mechatronics development. He is an author or co-author

of over 250 papers in journals and at international conferences and holds 14 patents. He has received the Best Jubilee Video Award at IROS (2012), the 2009 PE Publishing Award from the Journal of Systems and Control Engineering and prizes for

Best Paper at ICAR (2003) and the Best Student Paper Award at Robio (2013). He was also a Finalist for Best Entertainment Robots and Systems-20th Anniversary Award at IROS (2007) and finalist for the Best Manipulation paper at ICRA (2012), the Best Conference Paper at Humanoids (2012), the Best Student Papers at Robio (2013) and ICINCO (2014). He has been in the Program Committee of over 60 international conferences including ICRA, IROS, RSS, HUMANOIDS, BIOROB and ICAR, and he is a Technical Editor of IEEE/ASME Transactions in Mechatronics and on the Editorial Board of Robotics and Automation Letters. Since 2013, he is also serving as a Visiting Professor at the Centre for Robotics Research (CORE), Department of Informatics, Kings College University, London, UK.