

The Mechanics and Control of Leaning to Lift Heavy Objects with a Dynamically Stable Mobile Robot

Fabian Sonnleitner¹, Roberto Shu², and Ralph L. Hollis³

Abstract—A control algorithm is developed to enable dynamically stable spherical-wheel robots (ballbots) with arms to detect a heavy object of unknown mass, navigate to it, lift it, transport it, and place it in a desired location semi-autonomously. Previous work has successfully demonstrated two-wheeled dynamically stable mobile manipulator robots transporting heavy objects. We report here the first ballbot to reliably achieve such a task. A successful semi-autonomous lift and transport of a 15 kg heavy box whose actual mass was unknown was achieved using a combination of feedforward and feedback control laws based on a quasi-static center of mass computation. The ballbot's pan and tilt sensor turret tracked fiducial markers on the box. Ballbot-to-human and human-to-ballbot exchanges of a 10 kg heavy object was achieved while dynamically balancing.

I. INTRODUCTION

To be truly useful, mobile robots need to be able to lift and transport objects to collaborate with humans in work and home environments. This paper concerns the successful detection, lifting, and transport of a heavy object with a dynamically stable mobile robot. Heavy repetitive lifting continues to be a major health risk for workers. Further, current mobile manipulators are either too weak to carry heavy objects greater than about 10 kg or they have wide multi-wheel bases that make maneuvering difficult in narrow workplace or home environments. The ballbot is a radically different kind of mobile robot. Its large single spherical wheel gives it omnidirectional motion and omnidirectional compliance. Unlike traditional statically stable mobile robots which have wide bases to avoid tipping, gravity-referenced ballbots can be tall and slender with high centers of gravity. We show in this paper how ballbots can lift and transport heavy loads by *leaning back* to maintain stability, similar to how humans and humanoid robots carry out these tasks.

We employ the person-size CMU ballbot [1], described in more detail in Sec. II-C and shown in Fig. 1, for our study. We postulate that the ballbot's small footprint, strong arms and its ability to exert forces by inducing a lean angle make it a fitting robot to take over many demanding lifting tasks. To accomplish the lifting and transporting tasks autonomously, the ballbot must navigate to the location of the heavy object to be picked up (say, from a table), detect and recognize the object, grasp the object, pick up the object, transport the




Fig. 1. The CMU ballbot lifting a 15 kg payload using its 2-DOF arms while maintaining a fixed location on the floor.

object, and set the object down at a goal location. In our case, we relax the requirement of full autonomy for expedience but retain semi-autonomous operation. Alternatively, the ballbot can get and give the object from/to a human. If the human receives the heavy object from the ballbot, he/she can readily assess the object's weight by observing the lean angle of the ballbot. It is important that the ballbot retain its smooth and graceful balancing motion during all phases of the task, since people feel very uncomfortable around robots having high jerk motions [2]. To our knowledge, this is the first time that a ballbot is able to semi-autonomously pick/place heavy objects of unknown mass.

II. RELATED WORK

A. Human Lifting

As Hsiang explains in [3] there is a lot of research in human lifting motions and the resulting injuries and problems such as low back pain. According to the work of Hsiang, there are many studies on the risks of repetitive lifting, such as the study of Antwi-Afari [4]. But it is not fully understood how and if proper training and special lifting techniques help to avoid or reduce the risks. The study of Antwi-Afari performed a biomechanical analysis on the risks

¹ Fabian Sonnleitner is with The University of Applied Sciences, Salzburg, 5412 Puch bei Hallein, AUT sonnleitner.fabian@me.com

² Roberto Shu is with The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA rshu@andrew.cmu.edu

³ Ralph L. Hollis is with The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA rhollis@andrew.cmu.edu

for musculoskeletal disorders due to repetitive heavy lifting. They concluded that the main risk factors for musculoskeletal disorders are the lifting duration, weights, and postures.

Studies by Graham [5] and Bosch [6] determined that passive exoskeletons can reduce the muscle activity by about one third. Additionally, these devices can extend the endurance of the workers by almost three times. The system used in this study reduced the back issues but increased the discomfort in the chest region and is therefore not a good solution. These studies show that repetitive lifting tasks are bad for a person's well being, even with proper training and a good lifting position and additional assisting devices.

Hsiang explains in [3] that back injuries due to repetitive heavy lifting are a big cost factor for worker compensation costs in companies. The presented work could not find a representative correlation between different lifting techniques and low back pain. Therefore they conclude, that even if the safety officials of companies provide lifting training, the workers will still continue carrying heavy objects and therefore can suffer from the risks. This poses the question: "how can a person judge when an object is too heavy to safely lift?" A person detects the weight of the object either by judging from visual observation or by experimenting with the object. Already in the early childhood, humans learn to estimate the weight of an object just by observing another human interacting with it. If the person is not able to observe an interaction with the object, its mass is detected by the pressure applied to the fingertips and hands. Therefore, a human is able to quickly counterbalance the lifted object. But the quick counterbalancing reaction, however, can potentially injure muscles and tendons in the back.

Mechanized lifting devices have, of course, been around for eons. In robotic form there are autonomous fork lift trucks that can lift and transport massive payloads. A research question that remains is whether more agile machines such as humanoid robots can adequately lift and transport heavy loads in everyday human environments.

B. Dynamically stable and humanoid robots

Over the last decade, a variety of different humanoid robots were developed. Humanoids such as Atlas [7], Toro [8] and HRP-2 [9] showed successful lifting of payloads up to 12.2 kg. Other dynamically stable robots on two wheels such as Golem Krang and Handle were able to lift and carry oddly shape objects, such as a chair [10], [11], [12], and objects up to 45 kg [13]. (It should be recognized that two-wheeled mobile robots are dynamically stable in only one direction and are not omnidirectional.)

C. The CMU ballbot

The CMU ballbot is the first successful single-spherical-wheel mobile robot [14]. A description of the evolved robot can be found in [1]. The CMU ballbot is human size and can withstand hard and soft pushes by humans without losing its balance [15]. In 2011, a pair of 2-DOF arms with series-elastic actuators were added to the ballbot research platform and a shape space planner was developed which used the

arms to assist in navigation [16]. This planner enables arm movement with precisely known attached weights of up to 2 kg while the ballbot maintains a fixed position on the floor [1]. The ballbot is able to accomplish this by leaning backward so its center of gravity is directly above its center of support as illustrated in Fig. 2. This paper extends these capabilities to lift objects of much heavier unknown masses and transport them from place to place. To our knowledge, the ballbot is the first single-spherical-wheel robot capable of transporting heavy objects. Fig. 1 shows the CMU ballbot holding a 15 kg mass while balancing and maintaining a fixed position on the floor.




Fig. 2. Planar simplified ballbot model with arms and external payload. Note that the COM axis angle with respect to the vertical will change as arm angle α and payload mass m_{arm} change.

D. Differential Flatness Path Planning

The differential flatness path planner presented by Shomin [17] enables the ballbot to smoothly navigate autonomously through the environment while avoiding static and dynamic obstacles. The planner is capable of generating dynamical feasible trajectories for the underactuated ballbot online on the onboard computer. This was achieved by linearizing the system to find flat outputs, which is only possible if the ballbot's lean angle ϕ is assumed to be bounded between $-5^\circ < \phi < 5^\circ$, see Fig. 2. While ballbot's acceleration is proportional to its body lean angle, its position depends on the ball angle θ . To enable a point-to-point motion with an initial and final target condition, the *crackle* of the flat output is minimized, resulting in a ninth-degree polynomial trajectory for the ball path (θ_p) on the floor. Using the flat outputs and the 9th order polynomial a feasible lean angle trajectory (ϕ_p) can be generated [17]. In this work we modify the planner to be able to generate feasible trajectories when the initial state of the ballbot has a non-zero lean angle.

III. LEAN ANGLE COMPENSATION

The overall task consists of four separate blocks. The object detection and localization, the weight estimation, the center of mass calculation, and the lean angle compensation. The following paragraphs will provide deeper insights into these four blocks.

A. Object localization

The ballbot has a pan/tilt turret at its top that carries an RGB-D camera and other sensors. The camera was used to detect and locate the object to be picked up. Fiducial markers are widely used in robotics for detection, localization and mapping. To ensure a reliable object detection and localization the ARUCO framework presented in [18], [19] was used to avoid having to solve a more complicated vision problem. Due to the different positions of the object (a yellow box containing various weights and equipped with handles), on which a pair of 15×15 cm fiducial markers were affixed to ensure detection at a distance, as well as in close proximity. Fig. 3 presents the results of a benchmark test, where the target object to be lifted was placed on a table of 0.9 m height. During the test, the box was moved around randomly between the ballbot's front and 90° to the right side. The benchmark script was running at 5 Hz. Each time a valid pose of the object was detected it was marked on the map. The red crosses symbolize a detection of the top marker, the green crosses symbolize the location of the marker on the front side of the box. Because the object would interfere with the ballbot's body no detections were seen closer than 0.3 m. The top marker is detected up to a maximum of 2 m radius from the ballbot. The marker in the front could be detected up to 3.5 m away. Since the ballbot's pan/tilt turret can rotate 360° in pan, the result received from this 90° benchmark test is valid all around the ballbot.




Fig. 3. Results: Object detection and localization with ARUCO fiducial tracking.

B. Payload mass estimation

The weight of the box is detected by the torsion of the series elastic springs in both arms and the arm angle α . A polynomial function was generated for each arm to compensate the small differences between them. As each arm published its spring torsion and arm angles to a ROS topic, a separate laptop computer running Matlab™ can read and recorded this topic via Wi-Fi. Each arm was moved several times with different weights, commanded by a Matlab™ script. During those movements, the spring values, as well as the corresponding arm angles and attached weights, were recorded. Fig. 4 shows the curve fitting tool of Matlab™ which was used to fit a polynomial curve into the received data. The X and Y axis shows the arm angle and the spring torsion, while the Z axis shows the correlating mass. Several mass levels are visualized, as the mass was increased in 1 kg steps during the calibration. The procedure was done for each arm individually to achieve more precise results and to detect any load difference between the arms if the box is not centered during lifting.




Fig. 4. Determination of the polynomial function using Matlab's™ curve fitting tool.

C. Payload lean angle compensation

Lifting a heavy object away from the center of mass of the ballbot's body causes the entire system's center of mass to shift outside of the support point on the floor. This can lead to the ballbot falling over if the shift is not accounted for. We use the estimate of the mass m_{arm} and arm joint angle α to calculate the center of mass of the entire system. To calculate the new center of mass, the body center of mass position and the arm center of mass position are calculated separately as follows:

$$\begin{aligned} COM_{body,x} &= -L_b \cdot \sin(\phi) \\ COM_{body,y} &= L_b \cdot \cos(\phi) \end{aligned} \quad (1)$$

where $L_b = 0.71$ m is the body center of mass along the z axis and ϕ is the ballbot's lean angle.

$$\begin{aligned} COM_{arm,x} &= L_{arm,joint} \cdot \sin(\phi) - L_a \cdot \sin(\phi - \alpha) \\ COM_{arm,y} &= L_{arm,joint} \cdot \cos(\phi) - L_a \cdot \cos(\phi - \alpha) \end{aligned} \quad (2)$$

where $L_{arm,joint} = 1.3$ m is the shoulder joint height, $L_a = 0.56$ m is the arm length and α is the arm angle. Because

the lifting is either performed on the front or back side of the ballbot only a 2D model needs to be considered. For all the calculations we assume that the hollow thin aluminum tube arms are massless.

The body and arm center of mass offsets are combined to yield the new center of mass of the whole system as follows:

$$COM_{sys,x} = \frac{COM_{body,x} \cdot m_{body} + COM_{arm,x} \cdot m_{arm}}{m_{body} + m_{arm}}$$

$$COM_{sys,y} = \frac{COM_{body,y} \cdot m_{body} + COM_{arm,y} \cdot m_{arm}}{m_{body} + m_{arm}}$$
(3)

where $m_{body} = 81.65$ kg is the mass of the ballbot and m_{arm} is the mass of the object estimated using the procedure described in Section III-B.

Knowing the new system's center of mass, a desired body lean angle to bring the center of mass back on top of the support point is calculated by

$$\phi_a = \text{atan} \frac{COM_{sys,x}}{COM_{sys,y}}.$$
(4)

Combining Eq. 3 and Eq. 4 we numerically calculate the required body yield angle ϕ_a to maintain balance:

$$\phi_a = \text{acot} \frac{m_{body} \cdot L_b + m_{arm} \cdot (L_{arm,joint} - L_a \cdot \cos(\alpha))}{L_a \cdot m_{arm} \cdot \sin(\alpha)}$$
(5)

D. Lean angle control

The cascading control loop with feedforward compensation method used to lift and transport heavy objects is depicted in Fig. 5. The inner PID loop running at 500 Hz ensures that the ballbot maintains its balance at all times. It does so by tracking a desired lean angle, the details of which are described in [20]. In the case of purely balancing it tracks a zero lean angle. For other tasks such as lifting or navigation a non-zero lean angle is required. The inner loop is fed with three different feedforward compensation terms. The zero lean angle compensation term is a constant value measured offline to compensate for the body's center off mass being off of the vertical axis due to asymmetrical mass distributions inside the body. The simple PD feedback compensation is only active when the ballbot is in *station keeping mode*. This compensation prevents the ballbot's position from drifting if the zero lean angle compensation is not exactly correct. *Station keeping mode* is enabled right after the ballbot transitions from its statically stable state with all three legs deployed to its dynamically stable state with the legs retracted. The third compensation term corresponds to the new proposed strategy described in Section III-C. All of the compensation terms are fed to the inner loop at a rate of 500 Hz.

To autonomously navigate from point-to-point, station keeping mode is disengaged and the desired lean angle trajectory generated by the differentially flat path planner is fed to the inner loop. The path planner described in Section II-D was designed to operate between body lean

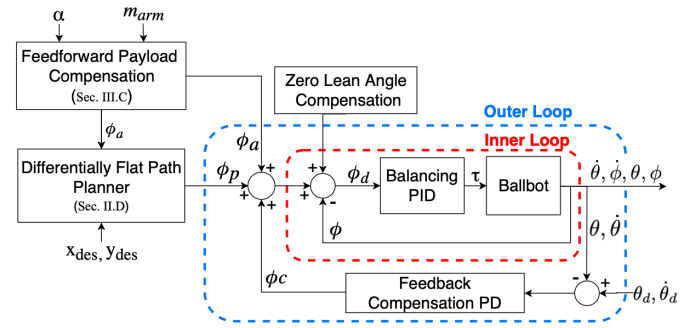


Fig. 5. Overview of the implemented cascading control loops with feedforward compensation terms.

angles ϕ of $-5^\circ < \phi < 5^\circ$. When the ballbot is carrying a heavy object there is an induced body lean angle ϕ_a needed to maintain balance. Without modification the path planner cannot account for the induced body lean angle and will generate trajectories that will destabilize the ballbot. Thus, we feed the induced body lean angle to the path planner. By tracking the desired body lean angle trajectory the inner PID loop will cause the ballbot to closely track a desired ground path [21].

IV. RESULTS

A. Lateral movement

Fig. 6 presents the lateral movement achieved while lifting a 10 kg heavy object with both arms. The actual lifting, where the arms were raised from 0° to 90° , occurred between second 5 and second 15. During the lifting maneuver, the object weight detection is constantly updated at a frequency of 500 Hz. This ensures that the object is immediately detected as soon as the “hands” touch it. During the whole maneuver, the unwanted ballbot wandering movement stayed below 2 cm, making it possible to do several experiments described in the following sections.

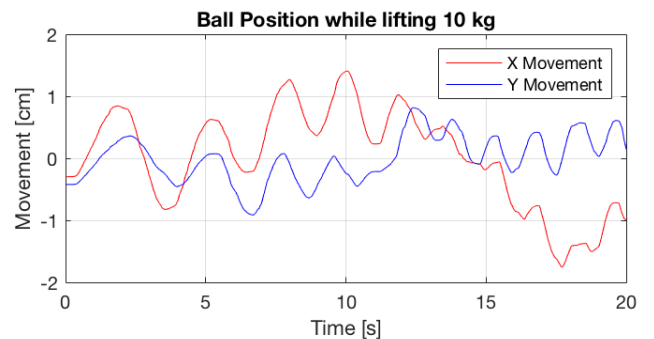


Fig. 6. Lateral movement while lifting a 10 kg heavy object

B. Human interaction

To be truly useful in a human environment, a ballbot needs to be able to deliver or receive objects at any time from a human. The following experiments present this ability.

The ballbot carries a 10 kg heavy object with its arms raised to 60°. The target is set 0.3 m in front of a subject person. Path planning is done by the differential flatness path planner as explained earlier in Sec. II-D. Fig. 7 shows that the target position near the human is reached. The subject grabs the object with both hands and lifts it within about 5 s. During this handover, the ballbot's position does not change noticeably, as the perceived load and lean angle is updated continually at 500 Hz. The load change on the arms is detected and the ballbot starts lowering its arms as soon as the payload reaches 0 kg. After the successful delivery, the ballbot drives back to its home position. This experiment presents the ballbot's ability to deliver an object with an unknown mass and adjust the lean angle fast enough to keep its position while interacting with a human. This test is performed with a maximum load of 10 kg.



Fig. 7. The ballbot delivers a 10 kg heavy object to a person

In a second experiment, the ballbot receives a 10 kg heavy object from a human subject. Again, the differential flatness path planner is used to drive close to the subject. After reaching the target position the subject places the object on the robot's extended "hands" at a normal human speed. The time between touching the robot's "hands" and the subject releasing his/her grasp was measured on average at 5 s. The ballbot immediately adjusts its lean angle to stay in the same position. Further, the ballbot informs the subject via speech output that it has detected an object and tells the subject how heavy it is. The ballbot then continues to drive back to a predefined location.

Both experiments were repeated four times to ensure repeatability. Each time the ballbot was able to counterbalance the additional forces and report the weight of the object within about ± 0.5 kg.

C. Transport

This experiment evaluates the ballbot's ability to semi-autonomously lift and transport an object between two different locations. The ballbot was positioned in front of the box prior to the start of the experiment. During the whole test, the ballbot was balancing gracefully. After sending the lifting command, the ballbot adjusts its arms slightly below the handles of the box and starts lifting the box with its pair of 2 DOF arms. The lifting of the box takes about 8 s. As

soon as the target arm angles of 80° is reached, the ballbot verbally reports the weight of the object and starts to yaw to its right side by 90°, as can be seen in Fig. 8.



Fig. 8. The ballbot is transporting a 10 kg heavy object from A to B.

After yawing to the correct position while transporting the 10 kg heavy object, the lowering of the object is initiated. The end position, with an object placed on the table, is reached 47 s after the lifting command is sent, as seen in Fig. 9. This experiment was repeated five times in a row with similar results. A similar experiment without performing the yaw motion was as well repeated five times successfully.



Fig. 9. The ballbot successfully placed the object on the table.

End-to-end autonomous transport between two remote locations was not repeatably achievable, as it was not possible to reliably navigate to pick up the object within the very tight margin of ± 4 cm needed for the "hands." Therefore, a successful lift was not guaranteed due to the imperfect starting position. The vision system detected the box reliably in 5 out of 5 trials, but a correct position which allowed a successful lifting was only reached twice. It is likely that this situation could be improved by introducing visual servoing techniques.

D. Heavy weight transport

This final experiment presents the maximum lifting and transport capability of the current setup of the ballbot. If there are tasks where higher limits would be needed, the series elastic actuators in the arms would need to be replaced or redesigned. During these tests, the spring loads were closely




Fig. 10. Closeup of the ballbot's arms and 3D-printed "hands" while holding the box and resting against the body.

monitored at all times to protect them from permanent deformation. For these tests, the arms were extended to 70° . Fig. 1 shows the ballbot carrying a 15 kg heavy object requiring the ballbot to lean back 6.2° . Note that after lifting a heavy object from a table and moving away, the arms can be lowered to bring the object into contact with the body before transporting it somewhere as shown in Fig. 10. This action, similar to the way a person might transport a heavy load, serves to reduce the lean angle and reduce or even eliminate power draw from the arms.

Even with a 15 kg heavy object the ballbot showed that it can keep its position stationary in station keep mode and yaw around its axis, at $10^\circ/s$, while keeping its position stationary within ± 2 cm. Additionally, the ballbot is able to navigate from point to point using the differential flatness path planning while carrying the 15 kg load.

V. CONCLUSIONS

This paper presented how a single wheeled dynamically stable mobile robot, such as the ballbot, is able to successfully lift and transport a heavy object. The results show for the first time that the ballbot class of mobile robots can lift and transport surprisingly heavy loads while retaining the ability to balance gracefully. We project that someday such robots can move out of the lab environment and into general use.

It was shown as well, through experimental results, that the CMU ballbot is now able to detect and locate objects using its onboard RGB-D camera and ARUCO fiducials. The lifting and carrying tasks showed that the ballbot was able to successfully lift and transport objects with varying weights of up to 15 kg. Additionally, it was shown that the ballbot can collaborate with humans by delivering or receiving heavy objects directly. The ballbot provides audio feedback to help the interacting humans to understand the ballbot's next movements and actions. The dynamic weight adaptation ensures that even if the weight changes suddenly, the robot will keep its balance.

The two main challenges were the correct weight estimation and the precision of the path planning. Due to the narrow opening of the "hands" and the relatively short limited-DOF arms, the position in front of the box needs to be as precise as ± 4 cm.

VI. FUTURE WORK

The ballbot relies on a body-mounted lidar to localize itself in mapped environments. While it transports heavy objects the stationary mounted lidar tilts at the same angle as the body. This tilt results in difficulties with correct localization. Therefore a different solution for the localization needs to be implemented. A possible solution could either be the same lidar mounted on a gimbal to compensate the lean angle or the use of a spherical lidar. Alternatively, visual SLAM and mapping could be implemented to improve the navigation. Also, the vision system needs to be able to detect objects without the need of markers.

Assuming the localization and object detection is solved, the next step is to grasp and lift the object. The current 2-DOF arms are restricted to very basic movements. The motion range can be visualized as two half spheres, each on one side of the ballbot. Therefore grasping smaller objects in front of the ballbot is currently not possible. The linear torsion springs limit the maximum weight of the objects carried to 15 kg. Additionally, due to the elastic elements on the arms joints, if the payload on the arms changes rapidly it will sometimes cause unstable vibrations at the arm joint. This phenomena can be observed in the accompanying video. In the short term a solution to this is to replace the current elastic elements with stiffer ones. In the long term our group is designing and building a pair of 7-DOF arms for the ballbot to overcome the limitations of the current arms.

The experimental results showed that the current path planner tracks long trajectories very well, but lacks in the precision of exactly reaching the target location. Because a precise position prior to lifting is essential, a more precise path planning method for short distances needs to be implemented. Additionally this new planner should consider the lean angle and hand position to navigate closer to the target object.

ACKNOWLEDGMENT

We want to thank the Austrian Marshall Plan Foundation for supporting this work and research period at the Robotics Institute of Carnegie Mellon University.

REFERENCES

- [1] U. Nagarajan and R. Hollis, "Shape space planner for shape-accelerated balancing mobile robots," *International Journal of Robotics Research*, vol. 32, no. 11, pp. 1323–1341, 2013.
- [2] B. Rohrer, S. Fasoli, H. I. Krebs, R. Hughes, B. Volpe, W. R. Frontera, J. Stein, and N. Hogan, "Movement smoothness changes during stroke recovery," *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, vol. 22, no. 18, pp. 8297–8304, 2002.
- [3] S. M. Hsiang, G. E. Brogmus, and T. K. Courtney, "Low back pain (LBP) and lifting technique - A review," *International Journal of Industrial Ergonomics*, vol. 19, no. 1, pp. 59–74, 1997.
- [4] M. F. Antwi-Afari, H. Li, D. J. Edwards, E. A. Pärn, J. Seo, and A. Y. Wong, "Biomechanical analysis of risk factors for work-related musculoskeletal disorders during repetitive lifting task in construction workers," *Automation in Construction*, vol. 83, no. June, pp. 41–47, 2017.

- [5] R. B. Graham, M. J. Agnew, and J. M. Stevenson, "Effectiveness of an on-body lifting aid at reducing low back physical demands during an automotive assembly task: Assessment of EMG response and user acceptability," *Applied Ergonomics*, vol. 40, no. 5, pp. 936–942, 2009.
- [6] T. Bosch, J. van Eck, K. Knitel, and M. de Looze, "The effects of a passive exoskeleton on muscle activity, discomfort and endurance time in forward bending work," *Applied Ergonomics*, vol. 54, pp. 212–217, 2016.
- [7] Boston Dynamics, "Getting some air Atlas," 2018. [Online]. Available: <https://youtu.be/vjSohj-Iclc> (Accessed 2018-06-01).
- [8] B. Henze, M. A. Roa, and C. Ott, "Passivity-based whole-body balancing for torque-controlled humanoid robots in multi-contact scenarios," *International Journal of Robotics Research*, vol. 35, no. 12, pp. 1522–1543, 2016.
- [9] H. Arisumi, J. R. Chardonnet, A. Kheddar, and K. Yokoi, "Dynamic lifting motion of humanoid robots," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 2661–2667, 2007.
- [10] Georgia Tech, "Golem Krang, The strongest humanoid robot," 2010. [Online]. Available: <https://youtu.be/7TjKxjQOzVw> (Accessed 2018-05-05).
- [11] M. Stilman, "Optimized Control Strategies for Wheeled Humanoids and Mobile Manipulators," *IEEE International Conference on Humanoid Robotics*, vol. December, 2009.
- [12] M. Stilman, J. Olson, and W. Gloss, "Golem Krang: Dynamically stable humanoid robot for mobile manipulation," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 3304–3309, 2010.
- [13] Boston Dynamics, "Introducing Handle," 2017. [Online]. Available: <https://youtu.be/-7xvqQeoA8c> (Accessed 2018-06-01).
- [14] T. Lauwers, G. Kantor, and R. Hollis, "One is enough!" *Proc. Int'l Symp. for Robotics Research*, vol. October, 2005.
- [15] U. Nagarajan, G. Kantor, and R. L. Hollis, "Human-robot physical interaction with dynamically stable mobile robots," *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction - HRI '09*, vol. March, pp. 281–282, 2009.
- [16] U. Nagarajan, "Fast and Graceful Balancing Mobile Robots," MS Thesis, Carnegie Mellon University, 2012.
- [17] M. Shomin and R. Hollis, "Differentially flat trajectory generation for a dynamically stable mobile robot," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 4467–4472, 2013.
- [18] S. Garrido-Jurado, R. Muñoz-Salinas, F. Madrid-Cuevas, and M. Marín-Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, pp. 2280–2292, 2014.
- [19] F. J. Romero-Ramirez, R. Muñoz-Salinas, and R. Medina-Carnicer, "Speeded up detection of squared fiducial markers," *Image and Vision Computing*, vol. 76, no. June, pp. 38–47, 2018.
- [20] U. Nagarajan, G. Kantor, and R. Hollis, "Integrated planning and control for graceful navigation of shape-accelerated underactuated balancing mobile robots," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 136–141.
- [21] M. Shomin and R. Hollis, "Fast, dynamic trajectory planning for a dynamically stable mobile robot," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*. IEEE, 2014, pp. 3636–3641.