Human Push-Recovery: Strategy Selection Based on Push Intensity Estimation

Lukas Kaul and Tamim Asfour, H^2T *

Abstract

We present methods for extracting a fast indicator for push-recovery strategy selection from the data of an inertial measurement unit mounted on a human or humanoid torso. The methods serve the purposes of detecting the beginning, the direction and intensity of the push as well as predicting the feasible recovery strategy. We test our methods on a dataset of 78 push-recovery trials collected with a human subject and show the relation between push intensity and recovery strategy.

1 INTRODUCTION

In the research field of humanoid robotics there are great efforts being made on addressing the problem of foot step generation, trajectory optimization and motion planning for bipedal walking (e.g. [1, 2, 6, 7]). These works demonstrate that keeping a two legged robot with human proportions stably balancing is a very difficult task, even under controlled conditions. However, if humanoid robots are to leave the research laboratories and fulfill tasks autonomously and reliably in the field, even more attention must be given to robust bipedal locomotion, the basis of such operation.

It became evident in the finals of the DARPA Robotics Challenge (DRC), held in June 2015, that even the most advanced humanoid robots do not yet show the degree of robustness in their locomotion capabilities that would be desirable and necessary to reliably fulfill outdoor missions [8]. Due to their technical complexity and fragility, most humanoid robots are at a high risk of damage from falling during walking, and many are not able to resume operation autonomously after such an incident even if undamaged (noteworthy exceptions being the very robust Atlas robot [9] and the CHIMP robot [10] that succesfully got up after falling during the DRC finals). This is why the research on robust balancing and walking, equipping robots with the ability to keep their balance even under challenging conditions, is crucially important for the advancement of humanoid robotics. Not surprisingly, this problem has gotten significant attention from researchers over the last years (e.g. [5, 11, 12, 13]).

Disturbances from the planned motion can have a number of different causes: They can result from an unstable or moving ground (e.g. when walking over gravel), from collision with an obstacle that was not incorporated in the motion plan (e.g. tripping) or from collisions at torso height (e.g. being pushed). The case of external disturbances in the form of pushes is the one that is of the highest interest for human/robot interaction, that is if humans and humanoid robots are to work together in an otherwise controlled environment. This will be the case for the most immediately envisioned applications of such robots, be it in personal care, rescue, or as factory co-workers. The reaction to pushes as a source of disturbance (i.e. push-recovery) is the part of balance recovery that has been given the most attention in the robotics community. Hyon et al. for example have presented a framework for balancing an recovery from light pushes aimed at compliant interaction with humans [14].

If an upright standing humanoid robot is pushed from the back, there are three high-level strategies that it can apply to keep its balance [3]: The *ankle strategy* shifts the center of pressure (COP) by applying torque to the ankle pitch joints. This shift leads to a torque applied by the ground reaction force (GRF) around the robot's center of mass (COM) that can be sufficient to recover from the disturbance and let the robot remain balanced in the case of relatively small pushes. Using the *hip strategy*, the robot applies a torque to its hip pitch joint, causing its upper body to bend. The reaction torque of this motion acts around the robot's COM and can be sufficient to recover from the disturbed state, even for pushes that can not be handled with the ankle strategy. Thirdly, if neither of the first two strategies is feasible, the robot needs to *take a step* to let the GRF at the new position generate a torque around its COM large enough to compensate for the push and to come to a still stand again.

The ankle strategy is the most straightforward to implement and therefore of limited interest to the research community. The hip strategy and several extensions that involve the movement of other body parts to generate torques that stabilize the robot are significantly more complex and have gotten more attention (e.g. [15, 16]). One especially

*The authors are with the Institute for Anthropomatics and Robotics, High Performance Humanoid Technologies Lab (H^2T) at the Karlsruhe Institute for Technology, Adenauerring 2, 76131 Karlsruhe, Germany; lukas.kaul@kit.edu; Fax: +49 721 608 44077

interesting aspect of these works is the importance of angular momentum control, largely inspired from human observation and by the fact that the whole-body angular momentum appears to be precisely controlled in the human gait cycle [17]. Push recovery by stepping has been researched as a dynamics problem, and using simplified linearized models of the humanoid body such as the Linear Inverted Pendulum and its derivatives ([4, 18, 19]), elegant formulations for the necessary step length, such as the capture point concept [5], have evolved.

In this work, we address the question what high-level push recovery strategy to use when presented with a specific push. For this purpose we divide the possible recovery strategies in two categories: those that do not involve a step and those that involve stepping. The human with its extremely sophisticated locomotion and balancing capabilities is an ideal source of information that can help answer such complex questions. Learning from human observation is a promising line of research that is also being explored in the literature for learning tasks from humans, especially for humanoid robots that share a similar geometry with humans (e.g. [20, 21]). We suggest that the human example as a source of information should be further explored, and that the answer to the question of push-recovery strategy selection can be learned from human observation.

Deciding on a strategy is difficult, as it has to happen in a very early stage of the push recovery in order for enough time to remain to successfully execute the recovery action. Furthermore, the direction of the push has to be known as early as possible so that the robot is able to react appropriately. We explore the possibilities of using only minimal sensory input to answer the questions about what type of strategy to use and in which direction the step is applied. The sensors we investigate are an external force sensor that captures the applied push force and a body mounted inertial measurement unit (IMU) that captures the linear acceleration of the torso.

2 METHODOLOGY

To analyze the push that is exerted onto the human subject during our data acquisition trials we investigate both an external force sensor and a body mounted IMU.

2.1 Sensing Tools

Force Measurement: We custom built a specific tool for measuring the push force exerted onto a human subject during push recovery experiments. The heart of the tool is a very sensitive piezo-resistive load cell that can measure forces of up to 1250N. The differential analog signal of the force sensor is amplified and digitized using an analog to digital converter (ADC) with 14 bit resolution. Periodic sensor sampling is controlled by a microcontroller which streams the data to a PC over a serial interface. As the electronics of the tool are battery powered and the serial downlink for data transmission can be established wirelessly, it offers great flexibility during experiments. The load cell

Figure 1 The force sensing glove, equipped with a piezo-resistive load cell, a microcontroller and a wireless downlink for data transmission. The leather cover is folded back to expose the sensor.

is embedded in a 3D-printed adapter that includes a front plate for pushing as well as a linear guiding mechanism. All of this is fitted to a boxing glove, which allows for very intuitive handling of the sensing system (Figure 1).

Acceleration Measurement: For acceleration measurement we use an IMU from Xsens that provides 3D linear acceleration and the 3D orientation of the sensor. During our experiments, we attach the IMU to the front of the subject's torso by strapping it around the chest (see Figure 2). As we are only interested in the torso's linear accelerations and not the gravity component of the acceleration, we perform a data processing step to remove the effects of gravity in our measurements: Given the global orientation of the IMU we transform the gravity vector into the sensor coordinate system (CS) and subtract its three components from the sensor acceleration readings. What remains is the gravity free acceleration vector a_{gf} with

$$a_{gf} = a_m - \mathbf{R}_W^S \begin{pmatrix} 0 \\ 0 \\ g \end{pmatrix} \tag{1}$$

where a_m is the measured acceleration in the sensor CS, $\mathbf{R}_W^S$ is the transformation matrix from the world CS to the sensor CS computed from the IMU orientation, and g is the magnitude of gravity (i.e. $9.81 m/s^2$).

Relation between push force and torso acceleration: During our push trials, we apply the pushing force at approximately the same height that the IMU is mounted on the torso. Given the proportionality $F = ma$ between the force F acting on a body of mass m and its acceleration a, it is not surprising to see that the measured push force and the body acceleration show very similar progressions. This is especially true for the first 300 ms of the push. After this period, the reactive movements of the human add additional acceleration to the torso, which results in a deviation between the applied force and the measured acceleration (see Figure 3). As we are interested only in the very first moments of the push and its features that let us predict whether a step should be taken or not, force and acceleration data both serve our needs.

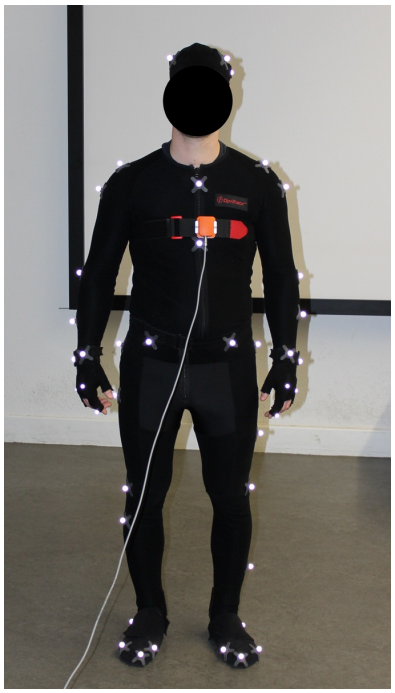

Figure 2 Subject wearing a motion capture suit and a chest mounted IMU for COM acceleration measurement. Note that motion capture data was not used for the presented study.

The IMU has two advantages over push force measurement: It not only provides a measure of the acceleration magnitude but also information about its direction, and it is much easier to include in a robot than force sensors that would allow measurement of interaction forces from various directions. For these two reasons we use the IMU sensor and the measured torso accelerations for quantifying a push trial rather than the push force itself, although both ways are feasible options.

2.2 Design of Experiments

We perform experimental push-recovery trials with a human subject in order to identify features in the torso acceleration that allow an early prediction of the type of push recovery strategy that will be used, and to identify the direction in which the push is applied. During the trials used for the presented study, the subject wearing the IMU is pushed from four different directions (right, left, back, front) and for each trial the applied strategy type (stepping, no stepping) is recorded. All pushes are applied at shoulder height either at the chest, the back or the shoulders.

3 EXPERIMENTAL RESULTS

3.1 Reference dataset

We collected a dataset that contains 78 push trials from the four directions back, front, left and right. Pushes were manually applied at different intensities so as to provoke different push-recovery strategies. Each trial consists of the

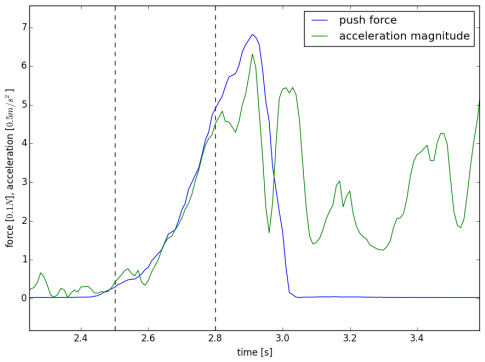

Figure 3 Push force and magnitude of linear torso acceleration during a relatively hard push from the back, scaled to match. Notice how the two are very similar in the first 300 ms of the push (between the dashed lines) and then diverge due to push recovery actions by the subject.

push force applied, measured with the force measurement glove, and the linear accelerations of the human subject, measured by the torso-mounted IMU. Both quantities are sampled at a frequency of 100 Hz and logged with consistent time stamps for synchronization. The dataset also contains a label for each trial, indicating the type of strategy performed by the subject, either *stepping* or *no stepping*.

At the beginning of each trial, the subject is advised to stand still with both feet placed in parallel and at shoulder-width. It is further advised to not to take a step for push-recovery unless it is necessary, in the sense that it can not otherwise recover from the push. This is the strategy that appears most feasible for robots: A step should only be taken if necessary, as it involves more power consumption, computational effort and a risk of falling by tripping or stepping onto unstable ground. These instructions aim at making the results applicable to humanoid robots and at making the experiments reproducible. However, a certain degree of unpredictability will remain due to the fact that there are many influences to the human behavior, like varying body tension, varying attention levels and unpredictable distractions.

From the data in the dataset we want to find ways to reliably detect a push, to compute the direction in which the push is applied and, most importantly, find a feature that allows an early prediction of the type of recovery strategy used by the subject. Learning rules to interpret this feature based on human observation will be useful for the development of efficient push-recovery algorithms for humanoid robots.

3.2 Push detection

The detection of the push in the continuous stream of relatively noisy IMU data is the first step in the data processing and the basis for all following steps. As it is not known initially in which direction the push was applied, the push detection routine should consider a quantity that is independent from this direction. As a direction-independent

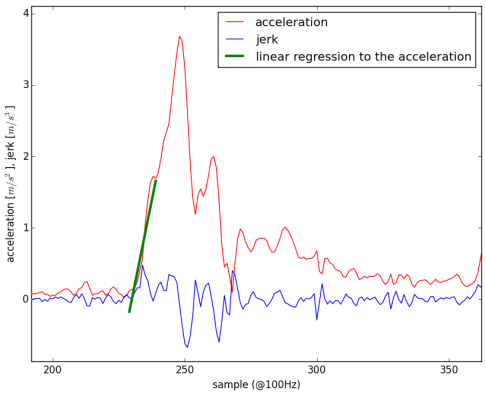

Figure 4 The measured magnitude of acceleration and the derived jerk during a push applied at the back. The green line shows the linear regression of the acceleration for the 10 first samples (i.e. 100 ms) after the push was detected.

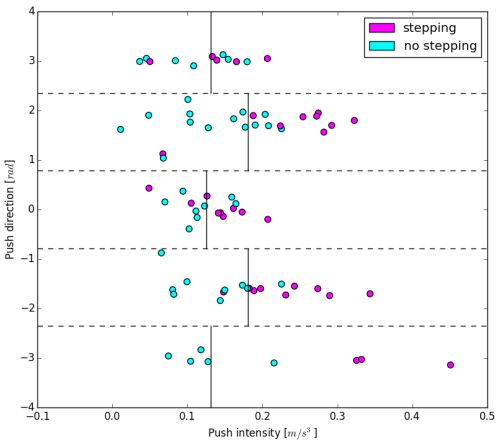

Figure 5 All 79 trials from the dataset, colored by the strategy observed during the experiments. The dashed lines indicate the separation by the estimated direction (dirctions of pi and $-pi$ both indicate a push from the back). The vertical lines are the borders that best separate the two labels in the three categories (push from back, front, side). The push intensity was estimated over the first 18 acceleration samples (180ms) of a push.

quantity we chose the scalar magnitude of the acceleration vector computed as

$$M_a = \sqrt{a_x^2 + a_y^2 + a_z^2} \tag{2}$$

where a_x, a_y and a_z are the gravity free components of the linear acceleration vector. To find the beginning of a push, a sliding window containing the newest and the last n_d acceleration samples is considered. A push is detected when the difference between the minimum and maximum magnitude value exceeds a specific threshold t_d. To add robustness against false positives to this method we require the maximum to occur after the the minimum, as a push in our experimental setup will always cause the torso acceleration to increase. We found that the start of any push in our dataset can be detected with values $n_d = 5$ and $t_d = 0.4m/s^2$. That means that a push can be detected based on the real-time IMU data as early as 50 ms after its beginning.

3.3 Push intensity

Our main intention is to identify and verify a feature that can be efficiently extracted from the IMU data and that allows for an early prediction of the push-recovery strategy that the human will chose in a specific trial. A decision strategy based on this feature derived from the human observation can than serve as a basis for efficient strategy selection on a robot. We investigate how we can use the intensity of the push as such a feature.

A direct measure for the push intensity would be the push peak fore or peak body acceleration. However, as can be seen in Figure 3, both the push force and body acceleration take a significant amount of time to build up. This is why the intensity of the push cannot be measured immediately after the push is detected. While one way of quantifying the push intensity would be waiting until a maximum acceleration is reached, this seems not to be a feasible approach. By the time the maximum acceleration is reached it might well be too late to properly react, even if it was clear then what the proper reaction to the push would have been.

Therefore, instead of the acceleration itself, we consider the the derivative of the acceleration magnitude with respect to time, or jerk, $\dot{M}_a$. Considering the time derivative seems feasible, since a steep increase in acceleration indicates a large maximum value, a strong push and therefore the execution of a stepping strategy. Similarly, a slow increase in acceleration indicates a weak push that might likely not require a step to be recovered from. The jerk itself, as the derivative of a measured quantity, is relatively noisy and therefore not straight forward to interpret (see Figure 4). To get a reliable estimation of the jerk value during the first moments of the push we approximate the jerk by fitting a least-squares linear regression to M_a over the first n_r samples after the start of the push. The slope of this regression line serves us as an easy to compute and interpret 1D feature that indicates the intensity of the applied push.

Upper and lower boundaries for the number of considered samples n_r have to be considered. A low number is desirable as it would allow to calculate the intensity feature shortly after the start, leaving more time for the execution of the push-recovery. However, with too few samples in consideration, the estimation is less robust against outliers in the data. With too many samples, the estimation of push intensity might become available too late, or acceleration samples that occurred after the peak are considered, which leads the linear regression to produce unreasonable results.

To evaluate if this jerk estimate is a valid feature for push recovery strategy prediction, we computed it for each trial in our dataset, varying the parameter n_r between 5 and 20

Table 1 Intensity based separation results for the pushes in our dataset, separated by direction.

Push direction	Correct strategy assignment
front	14 out of 18 (78%)
left and right	33 out of 40 (83%)
back	15 out of 20 (75%)

samples (i.e. using data between the detection of the push and the following 50 ms to 200 ms). We then find the border in terms of computed push intensity that separates the trials best into the two categories *stepping* (on the side of higher intensities) and *no stepping* (on the side of lower intensities). The best separation is the one that makes the smallest numbers of wrong assignments. Both intuition and our collected data indicate that push-recovery strategies vary with the direction in which the push occurs. Due to the feet being spaced at shoulder-width, larger pushes can be accepted without the need for stepping when they are applied from the sides than from either front or back, where the foot size limits the effective dimension of the support polygon. Taking into consideration these kinematic differences that cause different push-recovery strategy selection in the cases of the push being applied from the front, the back or the sides, we set three individual decision borders for those three cases. We found that for our reference dataset, the best separation results can be achieved with the number of acceleration samples considered for push intensity estimation n_r set to 18. With this number, decision borders that lead to smallest amount or wrongly predicted strategies are found. The decision boarders in this case are located at 1.25, 1.81 and $1.31 m/s^3$, for pushes from the front, the sides and the back, respectively. Table 1 summarizes the prediction results shown in Figure 5.

3.4 Push direction

For a robot to plan and execute a push recovery routine, it is essential to know the direction of the applied push. Only then is it possible to react in the right way, e.g. taking a step to the right when pushed from the left. Mathematically, computing the direction of a vector from its components is straightforward. Given the horizontal components a_x and a_y of the acceleration vector of a body, the horizontal direction of the acceleration can be computed with the two-argument inverse tangent function $atan2(a_x, a_y)$. The $atan2$ outputs an angle in the range from $-\pi$ to π. One way of making use of this is to compute the sample-wise acceleration direction using $atan2(a_x, a_y)$ for the same samples that are used to compute the push intensity as described in section 3.3. To reduce the dimension of the result to 1D, the mean value of the results is computed. This simple approach works well, except for push directions close to 180°, where $atan2$ has a discontinuity (switching from π to $-\pi$), causing the average to be somewhere in between,

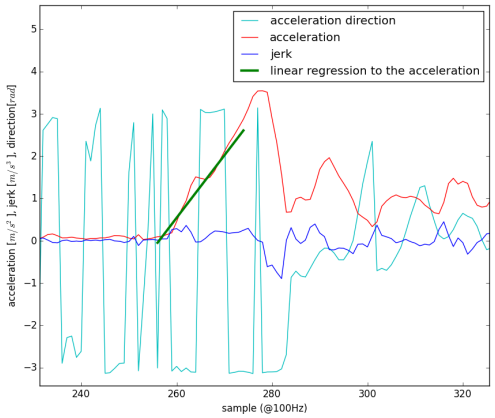

(a) Recorded data of a trial with a push applied from the back. Note how the sample-wise computed acceleration direction switches between π and $-\pi$ during the push.

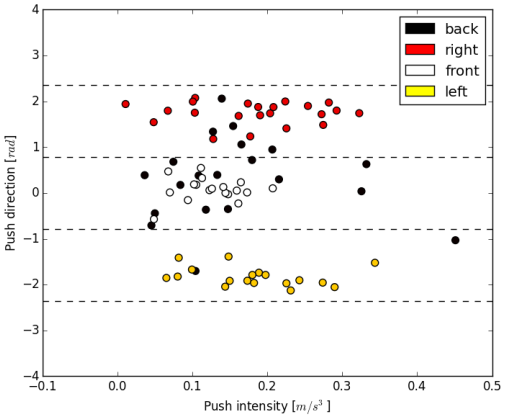

(b) The result of the direction estimate for 78 push trials using the mean value of sample-wise direction computation over the first 18 samples, colored by the actual direction. Pushes from the back cannot be correctly detected using this method.

Figure 6 Difficulties with detecting the push direction of pushes applied from the back when using the mean value of sample-wise computed acceleration direction.

but not where it should be (close to π or $-\pi$). This is illustrated with a sample push from the back and all pushes in the dataset in Figure 6.

To avoid this problem, we use a technique similar to the one used for push intensity estimation described earlier: We fit a least squares linear regression to the first n_r samples of the horizontal acceleration components and extract the regression line slopes $\dot{a}_{xr}$ and $\dot{a}_{yr}$. The push direction is then computed based on these two values as

$$d_p = atan2(\dot{a}_{xr}, \dot{a}_{yr}) \tag{3}$$

This method is robust against the discontinuity in $atan2()$ and correctly assigns the direction (π or $-\pi$) to all pushes from the back in the reference dataset. Only one push that was applied from the front is wrongly assigned a direction value in the 0.25π to 0.75π region, that indicates a push from the right. The results of the direction estimation for

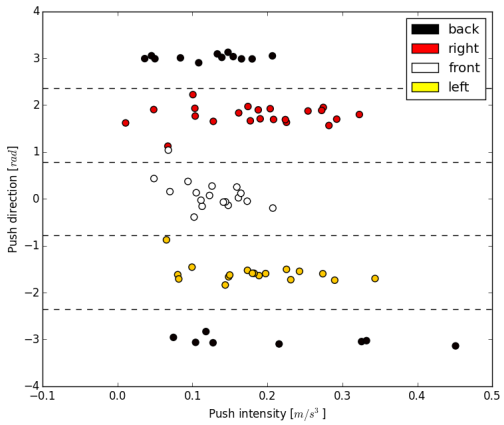

Figure 7 Results of push direction estimation using the slope of linear regression lines to the horizontal acceleration components as input. Note how pushes from the back are correctly classified.

all 78 pushes in our reference dataset are illustrated in Figure 7. Each datapoint represents one trial and is colored by the actual direction from where the push was applied. The dashed horizontal lines are spaced $\pi/2$ apart and mark the boundaries of the areas in which the pushes should lie (ideally at 0, $-\pi/2$, $\pi/2$ and $\pm\pi$ for pushes from the front, left, right and back, respectively).

4 CONCLUSION

4.1 Summary and discussion

In this study we aimed at finding ways of extracting useful information for bipded push recovery from a torso mounted IMU. We performed 78 push-recovery trials with a human subject in which we collected acceleration data and also recorded the type of strategy that the subject chose for push-recovery, either one that involves stepping or one that does not involve stepping. We aimed at extracting three pieces of information from the IMU data:

4.1.1 The point in time when the push begins

We compute the linear acceleration magnitude from the three acceleration components after we remove the influence of gravity. In a window that spans from the most recent acceleration sample back over a fixed number of older samples we compute the difference between the minimum and the the maximum acceleration magnitude value. If this difference is greater than a certain threshold, we conclude that a push started at the beginning of the inspected interval. We found that it is enough to take the 5 most recent IMU samples into consideration to reliably find the start of a push in our dataset. As we sample the IMU with a frequency of 100 Hz, this means that a push can be detected 50 ms after it began.

The evaluation of this push detection mechanism was performed only qualitatively, as there is no reference algorithm or ground truth to compare it against. We visually verified for all our trials that the linear regression line computed for intensity estimation starts at the beginning of the acceleration magnitude peak that indicates the applied push. An example of this can be seen in Figure 4.

The benefit of this push detection algorithm lies in its simplicity. Every computation involved in the push-recovery of a robot must be able to run in real-time, which motivates simple algorithms. The disadvantage lies in the two parameters, the number of considered samples and the detection threshold, that require experimental tuning.

4.1.2 The direction in which the push is applied

Estimating the direction of the applied push is equal to the problem of computing the direction of the torso acceleration. Two problems complicate this calculation: Firstly, the acceleration measurement is noisy so that some degree of smoothing or averaging needs to be applied. Secondly, when computing the direction from the horizontal acceleration components using the inverse tangent $atan2$ and averaging afterwards, the results are spurious for pushes applied in directions close to 180°due to the discontinuity in the inverse tangent. To get a reliable estimate of the push direction we fit linear regression lines over the first horizontal acceleration samples after the push was detected and compute the inverse tangent of the slopes of these lines. We computed this direction feature for every push in our dataset. For evaluation purposes we plotted the trials colored by the direction of the push along with lines that indicate the areas in the plot that indicate the four directions fron, right, left and back. Only one the direction estimate for one trial in our dataset lies in the wrong area of this plot (see Figure 7).

4.1.3 The intensity of the push

We hypothesized that the measure of intensity that we found can be an indicator for the push recovery strategy that the human subject will chose. Verifying this assumption was the main goal of our experiments. For this purpose we first defined a measure for the push intensity. We quantify the intensity of the push by the slope of the linear regression line to the magnitude of the torso's acceleration. This can be interpreted as a measure for the torso jerk. The jerk itself is relatively noisy and therefore difficult to work with, whereas the proposed estimate represents an easy to interpret 1D feature of the acceleration measurement during the push.

Our hypothesis was that this measure of push intensity early during the trial is an indicator for the type of push recovery strategy the human will chose. We found that, when considering the first 180 ms of acceleration data since the detected beginning of the push, the trials can be separated along the intensity-axis with a total of 79% trials being on the correct side - the pushes that provoked the subject to take a step on the side of higher intensity and the ones that did not require a step on the side of lower intensity (see Figure 5 and Table 1). To achieve this result we divided the

trials into three categories: those that involve steps from the back, the front an the sides. This seems feasible as the kinematics of the human and humanoid body require differing push recovery actions for pushes applied from different sides. We conclude from these results that it is possible to decide on a push recovery strategy based on the proposed push intensity estimation. We did not intend to train a general classifier as the exact value for intensity-based separation of pushes into the categories *requires stepping* and *does not require stepping* is not easy to generalize, especially not between different subjects/embodiments. However, we believe that the proposed measure can be a valuable component in push-recovery planning and execution, together with other sensory input available on a robot and that the observations we made will also be valid on humaonid robots, as their kinematics and dynamics are inherently similar to those of the human body.

4.2 Future work

In future work we plan to reproduce our experimental trials in the dynamic simulation of a full-size humanoid robot and further verify that push-recovery strategy selection based on the estimated push intensity from the torso acceleration during the first 180 ms of the push is a useful method.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Union Seventh Framework Programme under grant agreement no. 611909 (KoroiBot).

5 Literature

[1] R. Deits and R. Tedrake, "Footstep planning on uneven terrain with mixed-integer convex optimization", *IEEE-RAS International Conference on Humanoid Robots*, pp. 279–286, 2014

[2] J. Englsberger, T. Koolen, S. Bertrand, J. Pratt, C. Ott, and A. Albu-Schaeffer, "Trajectory generation for continuous leg forces during double support and heel-to-toe shift based on divergent component of motion", *IEEE International Conference on Intelligent Robots and Systems (IROS)*, pp. 4022–4029, 2014

[3] B. Stephens, "Push Recovery Control for Force-Controlled Humanoid Robots", Ph.D. dissertation, Carnegie Mellon University (CMU), 2011

[4] S. Kajita and K. Tani, "Study of Dynamic Biped Locomotion on Rugged Terrain - Derivation and Application of the Linear Inverted Pendulum Mode", *IEEE International Conference on Robotics and Automation*, pp. 1405–1411, 1991

[5] J. Pratt, J. Carff, S. Drakunov, and A. Goswami, "Capture Point: A Step toward Humanoid Push Recovery", *IEEE-RAS International Conference on Humanoid Robots*, pp. 200–207, 2006

[6] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point", *IEEE International Conference on Robotics and Automation*, pp. 1620-1626, 2003

[7] J. Chestnutt, M. Lau, G. Cheung, J. Kuffner, J. Hodgins and T. Kanade, "Footstep Planning for the Honda ASIMO Humanoid", *IEEE International Conference on Robotics and Automation (ICRA)*, 2005

[8] DARPA Robotics Challenge, `http://www.theroboticschallenge.org/`, Online, accessed: 2010-09-30

[9] ATLAS Robot, `http://www.bostondynamics.com/robot_Atlas.html`, Online, accessed: 2010-09-30

[10] A. Stentz, H. Herman, A. Kelly, E. Meyhofer, G. Clark Haynes, D. Stager, B. Zajac, J. Andrew Bagnell, J. Brindza, C. Dellin, M. George, J. Gonzalez-Mora, s. Hyde, M. Jones, M. Laverne, M. Likhachev, L. Lister, M. Powers, O. Ramos, J. Ray, D. Rice, J. Scheifflee, R. Sidki, S. Srinivasa, K. Strabala, J.-P. Tardif, J.-S. Valois, J.M. Vande Weghe, J. Michael, M. Michael and C. Wellington, "CHIMP, the CMU Highly Intelligent Mobile Platform", *Journal of Field Robotics*, volume 32, pp. 209-228, March 2015

[11] B. Stephens and C. Atkeson, "Push Recovery by stepping for humanoid robots with force controlled joints", *10th IEEE/RAS International Conference on Humanoid Robots*, pp. 52-59, 2010

[12] S.-J. Yi, B.-T. Zhang, D. Hong, D. D. Lee, "Online Learning of a Full Body Push Recovery Controller for Omnidirectional Walking", *IEEE-RAS International Conference on Humanoid Robots*, pp. 1-6, 2011

[13] J. Zhao, S. Schütz and K. Berns, "Biologically Motivated Push Recovery Strategies for a 3D Bipedal Robot Walking in Complex Environments", *IEEE International Conference on Robotics and Biomimetics*, pp. 1258-1263, 2013

[14] S.-H. Hyon, J. G. Hale and G. Cheng, "Full-Body Compliant Human – Humanoid Interaction : Balancing in the Presence of Unknown External Forces", *IEEE Transactions on Robotics*, volume 23, pp. 884-898, 2007

[15] S.-H. Lee and A. Goswami, "A momentum-based balance controller for humanoid robots on non-level and non-stationary ground", *Autonomous Robots*, volume 33, pp. 399-414, 2012

[16] A. Herzog, L. Righetti, F. Grimminger, P. Pastor and S. Schaal, "Balancing experiments on a torque-controlled humanoid with hierarchical inverse dynamics", *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014

[17] H. Herr and M. Popovic, "Angular momentum in human walking", *The Journal of experimental biology*, volume 211, pp. 467-481, 2008

[18] S. Kajita, F. Kanehiro, K. Kando, K. Yokoi, and H. Hirukawa, "The 3D Linear Inverted Pendulum Mode:

A simple modeling for a biped walking pattern generation", *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 239-246, 2001

[19] C. Doppmann, B. Ugurlu, M. Hamaya, T. Teramae, T. Noda and J. Morimoto, "Towards Balance Recovery Control for Lower Body Exoskeleton Robots with Variable Stiffness Actuators : Spring-Loaded Flywheel Model", *International Conference on Robotics and Automation*, pp. 5551-5556, 2015

[20] S. Nakaoka, A. Nakazawa, F. Kanehiro, K. Kaneko, M. Morisawa, H. Hirukawa and K. Ikeuchi, "Learning from Observation Paradigm: Leg Task Models for Enabling a Biped Humanoid Robot to Imitate Human Dances", *The International Journal of Robotics Research*, volume 26, pp. 829-844, 2007

[21] Y. Kuniyoshi, M. Inaba and H. Inoue, "Learning by watching: extracting reusable task knowledge from visual observation of human performance", *IEEE Transactions on Robotics and Automation*, volume 10, pp. 799 - 822, 1994