

Unsupervised Contact Learning for Humanoid Estimation and Control

Nicholas Rotella¹, Stefan Schaal^{1,2} and Ludovic Righetti^{2,3}

Abstract—This work presents a method for contact state estimation using fuzzy clustering to learn contact probability for full, six-dimensional humanoid contacts. The data required for training is solely from proprioceptive sensors - endeffector contact wrench sensors and inertial measurement units (IMUs) - and the method is completely unsupervised. The resulting cluster means are used to efficiently compute the probability of contact in each of the six endeffector degrees of freedom (DoFs) independently. This clustering-based contact probability estimator is validated in a kinematics-based base state estimator in a simulation environment with realistic added sensor noise for locomotion over rough, low-friction terrain on which the robot is subject to foot slip and rotation. The proposed base state estimator which utilizes these six DoF contact probability estimates is shown to perform considerably better than that which determines kinematic contact constraints purely based on measured normal force.

I. INTRODUCTION

Control and estimation approaches for legged robots rely on assumptions about the contact state of the feet. Floating-base inverse dynamics resolves underactuation by projecting the dynamics into the contact constraints, forcing the end-effector acceleration to be zero [1]. Locomotion on rough terrain focuses on stabilization through footstep adaptation but ignores the difficulties presented by contact constraint violations [2]. Similarly, legged odometry for base state estimation assumes the pose of an endeffector in contact is constant [3]. Methods have been introduced to robustify kinematics-based base state estimation, including computing a contact point with minimal instantaneous velocity [4] and outlier detection to discard measurements during slip [5], however few consider relaxing the contact assumptions by estimating contact quality in parallel.

Contact estimation is a broad topic which has been investigated in various contexts. Petrovskaya et al. [6] were among the first to consider multi-contact force control scenarios in which a manipulator interacts with the environment at points other than its endeffector. Del Prete et al. investigated the effect of contact point estimation error on force control for a humanoid with a capacitive skin [7]. Similar work was done by Manuelli et al. [8] to estimate contact points

without a tactile skin by fusing proprioceptive sensing with the dynamic model.

While estimation of contact points has been thoroughly studied, the problem of determining the *quality* of contacts is less well-defined. One aspect of contact quality is the determination of contact constraint directions for a given task. Ortenzi et al. [9] computed endeffector constraints for a manipulator in contact with a surface using only kinematics, and Nozawa et al. [10] recently presented a similar method for estimating environment constraints in humanoid control tasks.

For humanoids, the quality of a contact is largely determined by friction. Hoepflinger et al. [11] investigated foothold quality using unsupervised learning. Terrain elevation map samples were clustered to find a set of primitives which were evaluated for foothold robustness by computing the friction coefficient through exploratory force control. This allowed for prediction of contact quality from visual features for planning. While Focchi et al. [12] employed offline friction estimation through for a quadruped walking on steep slopes, Ridgewell et al. [13] introduced methods for online friction estimation and control adaptation. While the friction coefficient determines linear slip, the center of pressure (CoP) boundaries determine rotational slip/roll. Most controllers assume that the support polygon is the same shape as the foot, however this is invalid on rough terrain where line and even point contacts are encountered. Wiedebach et al. [14] presented one of the only approaches for online CoP boundary estimation during terrain exploration.

In contrast to approaches which indirectly compute contact quality by computing friction and CoP bounds, we wish to avoid contact models by directly estimating the probability of an endeffector being constrained in each of its six DoFs independently. In this direction, Hwangbo et al. [15] developed a Hidden Markov Model which uses kinematic and dynamic models to predict contact transitions without force sensing. This one-dimensional approach requires little sensing, however it does not estimate contact quality of the contact nor does it evaluate the classifier in an estimator.

Camurri et al. [16] recently developed a method for contact probability estimation using logistic regression. This one-dimensional classifier learns the normal force threshold at which the contact state transitions, ignoring lateral forces under the assumption that sufficient friction exists to prevent slip. The resulting probability per endeffector is used to weight the corresponding measurement in a base state estimator, and a heuristic for modulating the measurement variance to filter out the effect of impacts is introduced. Although this estimator performs better than one using a fixed threshold, the classifier requires significant effort to train -

This research was supported in part by National Science Foundation grants IIS-1205249, IIS-1017134, EECS-0926052, the Office of Naval Research, the Okawa Foundation, the Max-Planck-Society and the European Research Council under the European Union's Horizon 2020 research and innovation program (grant agreement No 637935). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding organizations.

¹Computational Learning and Motor Control Lab, University of Southern California, Los Angeles, California.

²Autonomous Motion Department, Max Planck Institute for Intelligent Systems, Tuebingen, Germany.

³New York University, New York, New York

ground truth is obtained manually as the contact sequence which minimizes estimation error. Further, all results shown are for walking on flat ground where slipping does not occur. Finally, only one dimension (normal force) is considered.

In contrast, we develop a contact estimator which:

- is completely unsupervised and model-free
- uses only common, proprioceptive sensors (endeffector force/torque and IMU)
- estimates the probability of contact in all six endeffector DoFs independently

We test this contact estimator for use in base state estimation by modulating the measurement uncertainty associated with each endeffector DoF using the corresponding estimated probability of contact. The following section details the motivation and setup for this approach.

II. BACKGROUND

A. Motivation

A difficult question arises when designing an estimator for contact state: what does it mean for an endeffector to be *in contact*? Most approaches treat the endeffector as fixed to contact surface if the normal force exceeds a chosen threshold, however contact truly occurs when the assumed endeffector constraints are satisfied - these statements are not always the same. The six DoF endeffector constraints are equivalent to enforcing that the feet cannot *slip* or *rotate*. An endeffector will not slip if the static friction constraint

$$\sqrt{F_x^2 + F_y^2} \leq \mu_{x,y} F_z \quad (1)$$

is satisfied, where F is the contact force and $\mu_{x,y}$ is the translational coefficient of friction. Likewise, the endeffector will not rotate if the CoP and rotational friction constraints

$$\begin{bmatrix} -\tau_y/F_z \\ \tau_x/F_z \end{bmatrix} \leq \begin{bmatrix} CoP_x \\ CoP_y \end{bmatrix} \quad (2)$$

$$|\tau_z| \leq \mu_z F_z \quad (3)$$

are satisfied, where τ is the contact torque, μ_z is the rotational coefficient of friction and CoP_x , CoP_y denote the contact support polygon bounds which are functions of contact surface geometry.

Since a sufficiently-high normal force F_z would guarantee that inequalities (1-3) are satisfied regardless of the other contact wrench dimensions, most estimation approaches simply threshold F_z [3], [17], [18]. However, this is restrictive especially on rough terrain where low friction and difficult surface geometry make slip and rotation likely even at high normal force values. It also results in a one-dimensional contact state estimate as in [15], [16], whereas the contact constraint is truly six-dimensional.

B. Sensing for Clustering

As discussed in the previous section, contact constraints are invalid when an endeffector slips and/or rotates, which is caused by a violation of friction and/or CoP constraints; these constraints depend on the contact wrench and surface properties (friction coefficients and geometry). Rather than estimate

these properties, we seek to cluster measured contact wrench data to directly learn constraint probabilities.

All experiments in this work are performed in the SL simulation environment [19]; we add simulated random-walk biases b_F and b_τ , along with simulated Gaussian noise processes w_F and w_τ , to the true force F and torque τ measurements:

$$F = \bar{F} + b_F + w_F \quad (4)$$

$$\tau = \bar{\tau} + b_\tau + w_\tau \quad (5)$$

As low-cost IMUs become available, humanoids are being augmented with additional sensing to improve estimation [20], [21]; in order to give structure to the clustering problem, we add a simulated IMU to each endeffector. We model the sensor outputs subject to simulated random-walk biases and thermal noise processes [22] as

$$a^{IMU} = R_W^{IMU} (a^W + g) + b_a + w_a \quad (6)$$

$$\omega^{IMU} = R_W^{IMU} \omega^W + b_\omega + w_\omega \quad (7)$$

where $a \in R^3$ and $\omega \in R^3$ are the linear acceleration and angular velocity, respectively. $R_W^{IMU} \in SO(3)$ is the rotation from world to IMU frame and g is gravity. Sensors are assumed to be aligned with the endeffector frame, however their positions relative to this frame origin are not required.

TABLE I: Simulated sensor noise standard deviations. Corresponding values for 1kHz sampling rate are shown.

	Continuous	Discrete (1kHz)
σ_θ	0.00000316rad/ \sqrt{Hz}	0.0001rad
σ_F	0.06325N/ \sqrt{Hz}	2N
σ_{b_F}	0.0001N/s/ \sqrt{Hz}	0.00316N/s
σ_τ	0.00316Nm/ \sqrt{Hz}	0.1Nm
σ_{b_τ}	0.0001Nm/s/ \sqrt{Hz}	0.00316Nm/s
σ_a	0.00078m/s ² / \sqrt{Hz}	0.02467m/s ²
σ_{b_a}	0.0001m/s ³ / \sqrt{Hz}	0.00316m/s ³
σ_ω	0.000523rad/s/ \sqrt{Hz}	0.01653rad/s
σ_{b_ω}	0.000618rad/s ² / \sqrt{Hz}	0.01954rad/s ²

III. CLUSTERING SETUP

Because we seek a continuous measure of contact quality rather than a classifier, we employ Fuzzy C-means (FCM) clustering which results in the soft partitioning of a dataset by allowing each data point to belong to more than one cluster [23]. This is accomplished by minimizing the cost

$$\sum_{i=1}^{N_p} \sum_{j=1}^{N_c} w_{i,j}^m \|x_i - c_j\|^2, \quad m > 1 \quad (8)$$

where N_p is the number of data points x_i , N_c is the chosen number of clusters, $w_{i,j}$ is the membership weight of point i belonging to cluster j and m is a constant which can be used to tune the amount of cluster overlap.

This cost is minimized in a manner similar to k-means clustering; first, initial membership weights are randomly assigned. Then, cluster means are computed as

$$c_j = \frac{\sum_{i=1}^{N_p} w_{i,j}^m x_i}{\sum_{i=1}^{N_p} w_{i,j}^m} \quad (9)$$

after which new membership weights are computed with

$$w_{i,j} = \frac{1}{\sum_{k=1}^{N_c} \left(\frac{\|x_i - c_j\|}{\|x_i - c_k\|} \right)^{\frac{2}{m-1}}} \quad (10)$$

Eq. (9-10) are iterated until the membership weights converge. Since $\sum_{j=1}^{N_c} w_{i,j} = 1$, we treat $w_{i,j}$ as the probability of point i belonging to cluster j .

We use an FCM implementation from the Python library Scikit-learn [24] with $N_c = 2$ clusters (corresponding to *contact* and *no contact* states) and default stopping parameters. The ‘‘fuzziness’’ constant is set to $m = 1.2$ which is the default value in most libraries. Increasing this factor can amplify the effect of slip on contact probability, however it also reduces the probability of contact when no slip occurs.

Each data point $x_k \in R^{7T}$ is a time series of the past $T = 20$ samples (at our control rate, $0.020s$). We include a short time-history to improve estimation response time; optimization of this time window is left to future work.

Clustering is performed independently for the six DoF $\{x, y, z, \alpha, \beta, \gamma\}$ of each endeffector, using the full contact wrench and the corresponding IMU dimension from $\{a_x^{IMU}, a_y^{IMU}, a_z^{IMU}, \omega_x^{IMU}, \omega_y^{IMU}, \omega_z^{IMU}\}$. The constraint in the local endeffector frame y direction uses, for example, data points of the form

$$x_k = \left\{ \{F_{x_{k-T}}, \dots, F_{x_k}\}, \{F_{y_{k-T}}, \dots, F_{y_k}\}, \{F_{z_{k-T}}, \dots, F_{z_k}\}, \right. \\ \left. \{\tau_{x_{k-T}}, \dots, \tau_{x_k}\}, \{\tau_{y_{k-T}}, \dots, \tau_{y_k}\}, \{\tau_{z_{k-T}}, \dots, \tau_{z_k}\}, \right. \\ \left. \{a_{y_{k-T}}^{IMU}, \dots, a_{y_k}^{IMU}\} \right\} \quad (11)$$

Data from sensors with noise added as in Sec. II-B is collected and used unfiltered for clustering. Preprocessing entails dimension-wise normalization of all x_k (to ensure that the scale of dimensions such as F_z do not dominate) followed by taking the absolute value (since slip is bi-directional).

IV. BASE STATE ESTIMATION

In order to evaluate the utility of the proposed contact estimator, we incorporate it into a base state estimation framework which relies on stationary contact assumptions. In previous work [25] we have implemented a kinematics-based estimator which fuses IMU data and relative base pose measurements to estimate the floating base state of a humanoid. The estimator measurements take the form

$$s_{p,i} = R(q)(p_i - r) + n_p \quad (12)$$

$$s_{z,i} = \exp(n_z) \otimes q \otimes z_i^{-1} \quad (13)$$

where $R(q)$ denotes the rotation matrix corresponding to the estimated base quaternion q , p_i and z_i are the estimated foot i position and quaternion respectively, and n_p and n_z are position and orientation measurement noise vectors (see [25]

for more details). In most approaches for legged robots, the variances of n_p and n_z are set to constant, tuned values and the measurements are dropped from the filter when the endeffector loses contact, determined based on a fixed normal force threshold.

In contrast, in this work we set the contact state (which determines active measurements) and measurement noise variance using the output of the probability estimator. When the probability of contact vector $P_{contact} \in R^6$ exceeds $P_i = 0.5$ in every dimension i , we consider the endeffector *in contact* and use the corresponding measurements. Further, we set the measurement noise covariance matrix as

$$\Sigma = E[nn^T] = r^2 I + \alpha(I - \text{diag}(P_{contact})) \quad (14)$$

where r is the nominal measurement noise standard deviation (sometimes tuned separately for position and orientation) and α is a scaling factor for the probability-dependent term (we choose $\alpha = 1$ for simplicity). The covariance thus converges to its constant value as in [25] when $P_{contact} \rightarrow 1$. This is conceptually similar to the approach of [16] but requires less tuning and considers all six contact dimensions.

Since clustering is performed in the endeffector frame, the covariance must be transformed into the base frame where the base state estimator measurement is expressed. This is accomplished with

$$\hat{\Sigma} = R \Sigma R^T, \quad R = \text{blockdiag}(R_{Endeff}^{Base}, R_{Endeff}^{Base}) \in R^{6 \times 6}$$

where $R_{Endeff}^{Base} \in R^{3 \times 3}$ is the rotation from endeffector to base frame (a function of kinematics and joint angles only).

V. EXPERIMENTS AND RESULTS

We perform a number of experiments to evaluate the performance of the proposed estimator and analyze its properties. All experiments are performed in the SL simulator [19] during a 60 second rough terrain walking task with simulated joint angle, IMU and contact wrench sensor noise as in Table I; noisy data is used for clustering, contact estimation and base state estimation. Control is computed using non-noisy sensor data and ideal base state estimation, however we investigate using the proposed contact estimator for closed-loop control in Sec. V-F. Walking velocity commands were recorded from user input and played back, producing repeatable trajectories across experiments. The rough terrain consists of raised patches with a friction coefficient of 0.4 (half the normal friction in SL). To account for the effect of noise and slight contact differences, Root-Mean-Squared Error (RMSE) for experiments in this section was computed by averaging performance across ten trials.

A. Contact Clustering Results

Sensor data was recorded from the rough terrain walking task and clustered offline as detailed in Sec. III (clustering takes on the order of a few seconds). The cluster means were then used to compute contact probability during a similar walking task; the results are shown in Fig. 1.

We focus on one contact cycle in the lower portion of Fig. 1 to investigate the clustering results more closely. Slip first

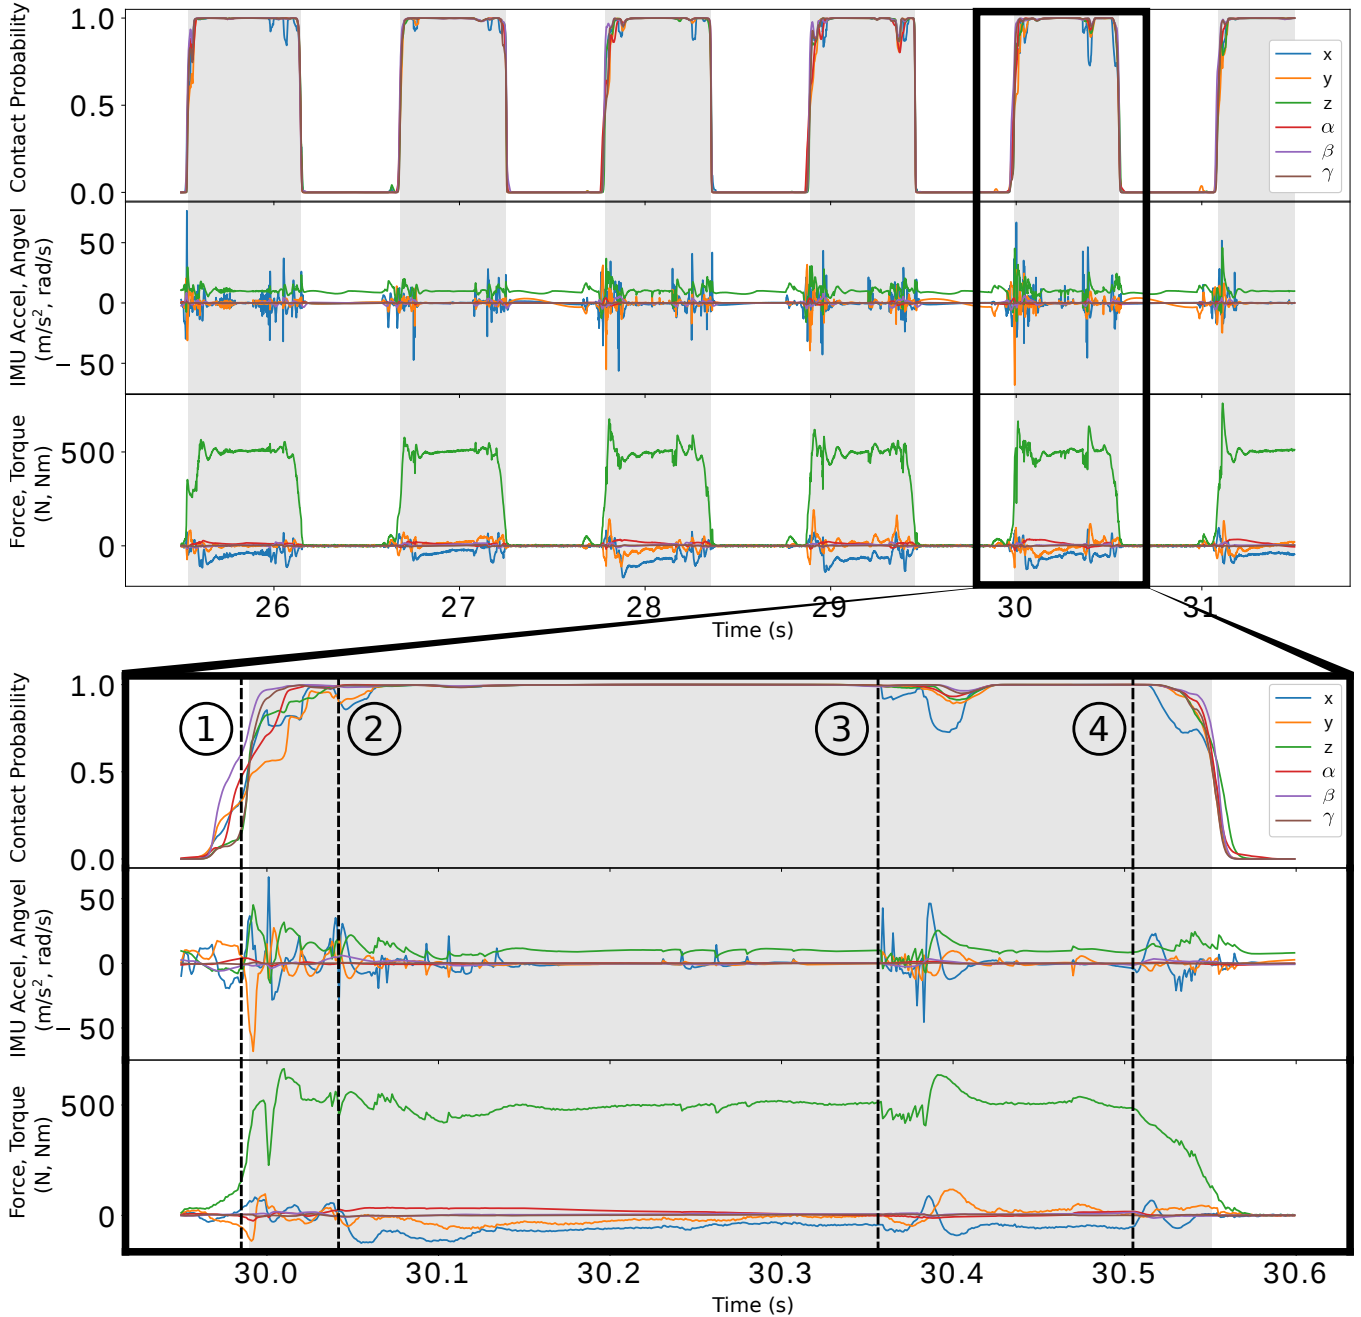


Fig. 1: The top portion shows the six-dimensional contact probability resulting from a rough terrain walking task (top) along with the measured IMU linear acceleration and angular velocity (middle) and measured contact force and torque (bottom). The portions in gray denote contact according to the probability estimator (all $P_i > 0.5$). The lower portion of the plot shows a zoomed view of one contact cycle with several distinctive contact events highlighted for discussion in Sec. V-A.

occurs in the y direction at (1), causing the corresponding contact probability to lag behind the other dimensions. Rotational slip in α is also present during loading, however on a smaller scale. Slip then occurs in x because the foot is not sufficiently loaded while the robot tries to create force in $-x$ to decelerate the center of mass; once F_z increases, there is sufficient friction to stop slipping and a negative F_x is sustained from (2) on. A drop in F_z during single support at (3) again causes slip, leading to a decrease in contact probability in all dimensions. Finally, slip in x again occurs

at (4) as the foot is being unloaded. These are only a few highlights of the complex contact interaction shown, however they aid in understanding where/why slip can occur.

B. Base State Estimation Threshold

Since we evaluate the proposed contact probability estimator against a typical humanoid base state estimator with a fixed normal force threshold for contact [25], we first perform experiments to optimize the chosen threshold. Performance is evaluated by computing the RMSE for the base position and

yaw angle as these four states are always unobservable without adding exteroceptive sensing. The normal force thresholds $\{10N, 40N, 100N, 200N, 400N\}$ were tested, with performance averaged across ten trials each; the results are shown in Fig. 2. The RMSE mostly decreased for increasing thresholds, with $200N$ resulting in the best performance. A threshold of $400N$ removes the double support period from estimation entirely, resulting in more error. We use a fixed threshold of $200N$ for the baseline estimator in experiments in the remainder of this work.




Fig. 2: Root Mean Squared Error (RMSE) for estimation of the unobservable base position (top) and yaw (bottom) for different normal force thresholds.

C. Clustering-Based State Estimation

We evaluate the base state estimator detailed in our previous work using both a fixed normal force threshold (as is commonly done) and using the proposed clustering-based contact probability estimator for the same rough terrain walking task. The base state estimators are identical other than the measurement noise covariance matrix modulation of Eq. (14). As shown in Fig. 3, using the contact clustering for base state estimation considerably reduces the RMSE.

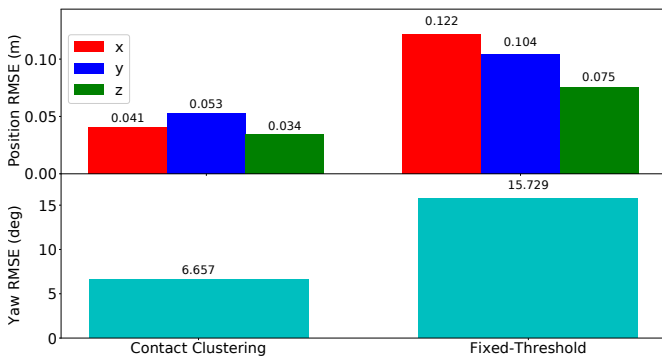


Fig. 3: Root Mean Squared Error (RMSE) for estimation of the unobservable base position (top) and yaw (bottom) for the contact probability-based base state estimator and the fixed normal force threshold base state estimator.

D. Clustering Training Data

In order to test how well the clustering-based estimator generalizes to different types of terrain, we performed clustering using data from two different tasks: one which walks over

rough terrain (as in all other experiments) and one which walks in place on flat ground. We then tested both contact estimators with separate base state estimators on the same rough terrain walking task; the resulting estimation errors are shown in Fig. 4.

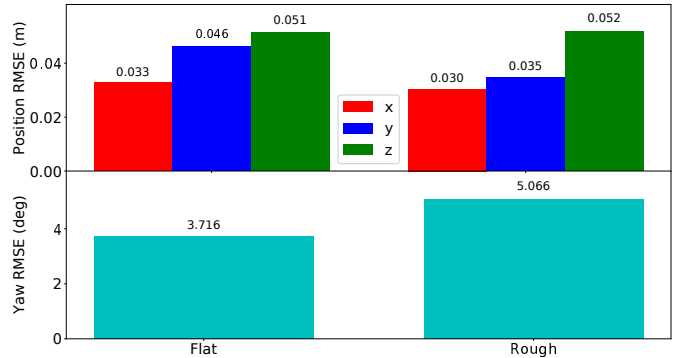


Fig. 4: Root Mean Squared Error (RMSE) for estimation of the unobservable base position (top) and yaw (bottom) for different training datasets.

Surprisingly, the estimator trained on flat ground walking data performs roughly equally-well, despite having been trained on a much different dataset than was used for testing. This is a desirable characteristic because obtaining data from rough terrain walking on a real robot is difficult, especially without accurate state estimation already in place.

We also wish to test how well the clustering-based estimator generalizes to different gaits. We perform clustering using data from flat ground walking in varying directions using three different gaits. The default gait used for walking in this work has a single support period of $0.5s$ and a double support period of $0.05s$; we denote this the *fast* gait. We also perform clustering on *slow* gait data (single support period of $1.0s$, double support period of $0.5s$). Finally, we cluster using data from a *mixed* gait which varies throughout the task between fast and slow. We then test the clustering-based estimators for these gaits for a mixed gait walk-in-place task on a path of rough terrain (varying the gait during a normal walking task over rough terrain is too unstable). The results are shown in Fig. 5.

The main conclusion which can be drawn from this study is that the best performance is obtained using the clustering trained on the mixed gait, as expected. However, the slow gait clustering generalizes much better than the fast gait clustering. The fixed-threshold base state estimator (denoted BSE) also performs quite well for this task, however because this was a walk-in-place there was mainly foot rotation and minimal slip; as seen from other tests, the clustering-based base state estimator performs much better when slip occurs. Further investigation into the effect of training data gait is left to future work.

E. IMUs for Clustering Versus Estimation

As motivated in Sec. II-B, the use of endeffector IMU data in addition to contact wrench data essentially supervises the clustering problem, since the accelerometer and gyroscope capture linear and rotational slip. We expect this sensor

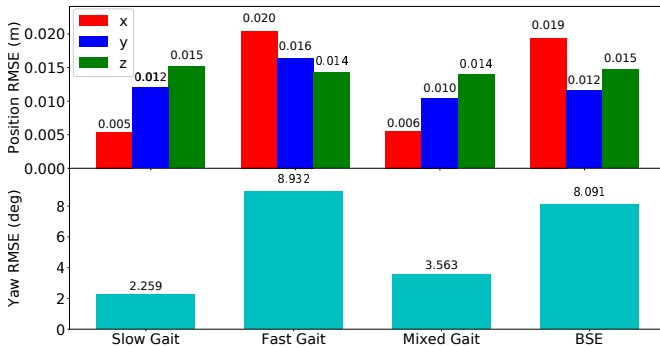


Fig. 5: Root Mean Squared Error (RMSE) for estimation of the unobservable base position (top) and yaw (bottom) for walking in place on a patch of rough terrain with a varying gait using clustering trained on three different gait types as well as for the fixed-threshold base state estimator (BSE).

data to embed structure in the resulting clusters, meaning that IMUs should not be required when running the contact estimator afterwards. To test this, we cluster using data points as in Eq. (11) but perform clustering-based state estimation with both a) the full data points including IMUs and b) without IMUs (dropping the last portion of Eq. (11)). The resulting estimation errors are shown in Fig. 6.




Fig. 6: Root Mean Squared Error (RMSE) for estimation of the unobservable base position (top) and yaw (bottom) with and without using IMU data for online contact estimation.

Although the RMSE is slightly lower in all dimensions when using the IMU data, performance is not considerably changed when it is removed. This is a very useful property because it means that the endeffector IMUs can be removed after initially collecting data for clustering. While some robots are designed with endeffector IMUs, most are not; using this clustering method would involve temporarily attaching IMUs as in [20], [21]. This is reasonable for training, however attaching these sensors permanently involves designing rigid mounts, routing cables and protecting them from collisions with the environment. The ability to remove IMUs after training the estimator is highly advantageous when working with real hardware.

F. Clustering-Based Contact Estimation for Control

While the primary focus of this work has been on the development and evaluation of a contact probability estimator for use in base state estimation, the proposed method has applications in humanoid control as well. The most direct

application is the use of an improved base state estimator in a walking controller such as the one used to generate data in this work.

This walking controller uses a simplified model in a model predictive control framework to plan center of mass and endeffector trajectories, which are tracked using an optimization-based inverse dynamics controller similar to [26]. The estimated base pose is crucial in computing both dynamic model parameters and feedback control for endeffector tracking. In the attached video, we demonstrate that the use of our contact probability estimator in this context improves control considerably, allowing the robot to walk on the rough terrain for a longer time before falling due to an accumulation of base state estimation error.

VI. CONCLUSIONS

The clustering-based contact probability estimator presented in this work estimates the quality of contact using only proprioceptive sensor data in a completely unsupervised approach. Unlike previous works, this estimator provides the probability of satisfying endeffector contact constraints in all six dimensions independently. Use of this method in a base state estimation framework was shown to considerably lower estimation error as compared to a base state estimator which uses a fixed normal force threshold and noise parameters. The proposed method also exhibits favorable properties which allow it to be used without endeffector IMUs after training and generalize to new terrain. Finally, it was shown that use of this improved base state estimator for closed-loop inverse dynamics control allows the robot to remain stable during rough terrain walking for longer. Future work will include further analysis of the properties of this contact estimator as well as a more low-level control application in which endeffector constraints in inverse dynamics are smoothly varied according to the contact probability.

REFERENCES

- [1] M. Mistry, J. Buchli, and S. Schaal, "Inverse dynamics control of floating base systems using orthogonal decomposition," in *2010 IEEE International Conference on Robotics and Automation*, May 2010, pp. 3406–3412.
- [2] V. Barasuol, J. Buchli, C. Semini, M. Frigerio, E. R. D. Pieri, and D. G. Caldwell, "A reactive controller framework for quadrupedal locomotion on challenging terrain," in *2013 IEEE International Conference on Robotics and Automation*, May 2013, pp. 2554–2561.
- [3] M. Bloesch, M. Hutter, M. Hoepflinger, S. Leutenegger, C. Gehring, C. D. Remy, and R. Siegwart, "State estimation for legged robots - consistent fusion of leg kinematics and IMU," in *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2012.
- [4] K. Masuya and T. Sugihara, "Dead reckoning of biped robots with estimated contact points based on the minimum velocity criterion," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 3637–3642.
- [5] M. Bloesch, C. Gehring, P. Fankhauser, M. Hutter, M. A. Hoepflinger, and R. Siegwart, "State estimation for legged robots on unstable and slippery terrain," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 6058–6064.
- [6] A. Petrovskaya, J. Park, and O. Khatib, "Probabilistic estimation of whole body contacts for multi-contact robot control," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, April 2007, pp. 568–573.
- [7] A. D. Prete, F. Nori, G. Metta, and L. Natale, "Control of contact forces: The role of tactile feedback for contact localization," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2012, pp. 4048–4053.

- [8] L. Manuelli and R. Tedrake, "Localizing external contact using proprioceptive sensors: The contact particle filter," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 5062–5069.
- [9] V. Ortenzi, H. C. Lin, M. Azad, R. Stolkin, J. A. Kuo, and M. Mistry, "Kinematics-based estimation of contact constraints using only proprioception," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, Nov 2016, pp. 1304–1311.
- [10] S. Nozawa, S. Noda, M. Murooka, K. Okada, and M. Inaba, "Online estimation of object-environment constraints for planning of humanoid motion on a movable object," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, May 2017, pp. 1291–1298.
- [11] M. A. Hoepflinger, M. Hutter, C. Gehring, M. Bloesch, and R. Siegwart, "Unsupervised identification and prediction of foothold robustness," in *2013 IEEE International Conference on Robotics and Automation*, May 2013, pp. 3293–3298.
- [12] M. Focchi, A. del Prete, I. Havoutis, R. Featherstone, D. G. Caldwell, and C. Semini, "High-slope terrain locomotion for torque-controlled quadruped robots," *Autonomous Robots*, vol. 41, no. 1, pp. 259–272, Jan 2017. [Online]. Available: <https://doi.org/10.1007/s10514-016-9573-1>
- [13] C. Ridgewell, "Humanoid robot friction estimation in multi-contact scenarios," Ph.D. dissertation, Virginia Polytechnic Institute and State University, 2017.
- [14] G. Wiedebach, S. Bertrand, T. Wu, L. Fiorio, S. McCrory, R. Griffin, F. Nori, and J. Pratt, "Walking on partial footholds including line contacts with the humanoid robot atlas," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, Nov 2016, pp. 1312–1319.
- [15] J. Hwangbo, C. D. Bellicoso, P. Fankhauser, and M. Hutter, "Probabilistic foot contact estimation by fusing information from dynamics and differential/forward kinematics," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 3872–3878.
- [16] M. Camurri, M. Fallon, S. Bazeille, A. Radulescu, V. Barasuol, D. G. Caldwell, and C. Semini, "Probabilistic contact estimation and impact detection for state estimation of quadruped robots," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 1023–1030, April 2017.
- [17] M. F. Fallon, M. Antone, N. Roy, and S. Teller, "Drift-free humanoid state estimation fusing kinematic, inertial and lidar sensing," in *2014 IEEE-RAS International Conference on Humanoid Robots*, Nov 2014, pp. 112–119.
- [18] S. Faraji, L. Colasanto, and A. J. Ijspeert, "Practical considerations in using inverse dynamics on a humanoid robot: Torque tracking, sensor fusion and cartesian control laws," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, Sept 2015, pp. 1619–1626.
- [19] S. Schaal, "The sl simulation and real-time control software package," University of Southern California, Dept. of Comp. Sci., Tech. Rep., June 2007.
- [20] N. Rotella, S. Mason, S. Schaal, and L. Righetti, "Inertial sensor-based humanoid joint state estimation," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1825–1831.
- [21] X. Xinjilefu, S. Feng, and C. G. Atkeson, "A distributed mems gyro network for joint velocity estimation," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016, pp. 1879–1884.
- [22] O. J. Woodman, "An introduction to inertial navigation," University of Cambridge, Computer Laboratory, Tech. Rep. UCAM-CL-TR-696, Aug. 2007. [Online]. Available: <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-696.pdf>
- [23] J. C. Dunn, "Well-separated clusters and optimal fuzzy partitions," *Journal of Cybernetics*, vol. 4, no. 1, pp. 95–104, 1974. [Online]. Available: <http://dx.doi.org/10.1080/01969727408546059>
- [24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [25] N. Rotella, M. Bloesch, L. Righetti, and S. Schaal, "State estimation for a humanoid robot," in *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, Sept 2014, pp. 952–958.
- [26] A. Herzog, N. Rotella, S. Mason, F. Grimmering, S. Schaal, and L. Righetti, "Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid," *CoRR*, vol. abs/1410.7284, 2014.