

Automated Acquisition of Anisotropic Friction

Keno Dreßel*, Kenny Erleben[†], Paul Kry[‡] and Sheldon Andrews[§]

*Department of Nordic Research, University of Copenhagen
Copenhagen, Denmark, Email: fwn391@alumni.ku.dk

[†]Department of Computer Science, University of Copenhagen
Copenhagen, Denmark, Email: kenny@di.ku.dk

[‡]School of Computer Science, McGill University
Montréal, Canada, Email: kry@cs.mcgill.ca

[§]Department of Software and IT Engineering
École de technologie supérieure de Montréal, Canada
Email: sheldon.andrews@etsmtl.ca

Abstract—Automated acquisition of friction data is an interesting approach to more successfully bridge the reality gap in simulation than conventional mathematical models. To advance this area of research, we present a novel inexpensive computer vision platform as a solution for collecting and processing friction data, and we make available the open source software and data sets collected with our vision robotic approach. This paper is focused on gathering data on anisotropic static friction behavior as this is ideal for inexpensive vision approach we propose. The data set and experimental setup provide a solid foundation for a wider robotics simulation community to conduct their own experiments.

Index Terms—friction; static friction; computer vision; automated friction measurements; robot arm;

I. INTRODUCTION

Friction models approximate the frictional behavior of surfaces in contact. Parameter tuning for these models is therefore vital for achieving simulation results that predict real-world contact dynamics [1]. However, roboticists and simulation experts have noted the reality gap introduced by the models used by off-the-shelf physics engines [2], [3]. It is a significant challenge in robotics to overcome this gap, and there have been attempts to reduce its impact or remove it all together [4], [5], [6].

The drawbacks of using approximate models are especially relevant when dealing with anisotropic friction, where the frictional forces depend on the relative orientation of the surfaces. In the case of Coulomb friction, which is the model we primarily address in this work, anisotropy is realized by two coefficients and results in an elliptical limit curve [7] rather than a circular one. While this covers some cases, there is a long list of possible limit curves that are not necessarily elliptical, or even convex [8]. Recent work has noted the randomness of motion due to frictional effects and proposed using a stochastic process to model friction [3].

We propose a system for measuring the static friction between two surfaces in which a platform is fixed to the end of a robotic arm, and a test object rests on top of the platform, but is otherwise free to move (see Figure 1). The platform is tilted to gradually increase the angle of incline, and a tracking system reports the relative movement of the test object. A large




Fig. 1: The robot arm tilts the platform until sliding motion is detected. This automated approach can measure static friction freely in many directions and with minimal bias.

database of movement with frictional effects is gathered for different test objects.

As the testing object is not pushed, pulled, or otherwise constrained, a natural sliding behavior can be observed. This minimizes the potential bias introduced by interactions with the measurement equipment and allows for the creation of a novel, rich, and detailed data set for a selection of common material examples. Our dataset is available online, is open, and includes open source code for data processing. To our knowledge our approach is the first using a fully automated passive strategy, rather than actively pushing or pulling objects as done in other work [3], [1].

The agile and inexpensive vision robotics approach make our paradigm easily accessible to others and is an important first step in feeding data into non-conventional models for friction.

II. RELATED WORK

Many models of friction have been developed. A good recent survey was done by Marques et al. [9]. Although there is a wide range of models available, a very commonly used model is still the isotropic planar Coulomb friction model. This model is limited to rough, dry and stiff surfaces [10] but

is often used in simulators as a default model for all types of surfaces. All other models have their own limitations as well. There is no universal model that can cover all variants of surfaces and modifiers at once.

For accurate, high fidelity physics simulations, detailed information about the friction must be available for the given surface pair. A common method is spring based static friction measurements. This involves using sensor in each direction that forces are measured. Other methods include pushing or pulling the object. More specialized setups are used for testing combinations, such as human fingers on surfaces [11], or tires on roads [12]. A good survey covering different methods has been compiled by Persson [13]. A common limitation is the single testing direction, unlike our approach which measures frictional behavior in all directions. Previous approaches require measurement experiments to be manually reset [14], which is an obstacle for gathering large amounts of data in an automated way. In contrast, our system is automated and does not need to be manually reset.

The setup presented in this work aims to take a modern approach where robots are used to circumvent both the limitations in testing directions and quantity of data. This concept can be successfully employed in friction research as shown by Pai et al. [15] and Yu et al. [3]. Additional advantages are high accuracy during the experiment on repeated motions and transparent testing procedures because the code is made available. Our approach uses a passive strategy for measuring friction and is focused on collecting anisotropic data.

III. THE ROBOTIC MEASUREMENT SYSTEM

Our measurement system consists of three main parts: the robotic arm, the platform mount, and the motion tracking system.

A. Robotic arm

The robotic arm we use in our experiments is a Kinova JACO 6-DOF prosthetic arm with ROS integration. The default Trac-IK inverse kinematics algorithm was used to navigate the robot in its workspace. A GitHub repository with source code and documentation of the ROS package we developed for controlling the robotic arm.¹

We use the internal sensors of the robot to measure the angle of inclination of the platform. To validate that this choice is reasonable we first confirm that the absolute orientation of the platform could be reliably recovered from the internal sensors using forward kinematics. We therefore performed a comparison of the orientation reported by the robot arm with other measurement devices. Specifically, a 12 camera OptiTrack motion capture system, a Capri 20005 Digital Angle Gauge, and inertial measurement unit (IMU) from a Pixhawk flight controller. An example showing the optical motion capture system and robot arm is shown in Figure 2c.

The comparison involved tilting the end-effector between 0 and 10 degrees five times. Figure 3 shows the difference




Fig. 2: At left (a), the gray 3D printed base has dark gray foam camera mount placed at a fixed offset with thin wooden posts. On the wooden test plate is a piece of white foam with a circular cutout, which prevents objects from leaving the test area. At top right (b), dark blue tape on the object allows us to track position and orientation of the object. At bottom right (c), the calibration setup used to determine the accuracy of the robotic arm’s self-reported angular orientations.




Fig. 3: Difference between reference and measured angles of the end-effector. The IMU has the most noise, and the motion capture appears to contain bias, but otherwise, the sensors follow the same pattern to within half a degree, suggesting that the internal sensors are a reliable data source.

of each measurement device versus a reference angle. Large errors are observed for the IMU, but there is a strong correlation between the other sensors. We observed that the motion capture had an offset of around 0.5 degrees, but otherwise the measured angles were in line with the inclinometer and the internal sensors. Therefore, we are satisfied that the internal sensors provide sufficient accuracy for this experiment. A full comparison of the different measurement devices can be seen in Table I.

¹<https://github.com/kenodressel/afm>

TABLE I: Variations in the angle of inclination measurements using different devices. All numbers are given in degrees. Each cell shows the mean absolute value of the difference between the two compared devices. Low values of the standard deviation indicate that the pair devices measure similar angles.

	Reference	IMU	Motion capture	Internal sensors	Inclinometer
Reference	0 ± 0	0.39 ± 0.42	0.60 ± 0.23	0.16 ± 0.25	0.21 ± 0.25
IMU	0.39 ± 0.42	0 ± 0	0.82 ± 0.39	0.42 ± 0.40	0.37 ± 0.38
Motion capture	0.60 ± 0.22	0.82 ± 0.39	0 ± 0	0.55 ± 0.05	0.65 ± 0.06
Internal sensors	0.16 ± 0.25	0.42 ± 0.40	0.55 ± 0.05	0 ± 0	0.10 ± 0.04
Inclinometer	0.21 ± 0.25	0.37 ± 0.38	0.65 ± 0.06	0.10 ± 0.04	0 ± 0




Fig. 4: The center of rotation was matched to the center of mass from an even $1 \text{ cm} \times 1 \text{ cm} \times 1 \text{ cm}$ object resting on the base plate (b) in order to minimize side-effect of any linear acceleration of the object caused by the robot (a).

B. Platform mount

The standard end-effector of the arm was replaced with a custom 3D printed end-effector on which we mount our test surfaces. We subsequently adjusted the center of rotation in the unified robot description format (URDF) file in order to minimizing the linear acceleration of samples placed on the test plate. Without this adjustment, the test object could otherwise experience premature sliding, as illustrated in Figure 4.

The 3D printed mount can fit any surface that can be cut to a $0.1 \text{ m} \times 0.1 \text{ m}$ square. The foam border in Figure 2a stops the object from sliding into corners, as this would produce additional border contact while sliding. The soft foam permits both easy rapid prototyping and withstands impacts without any permanent deformations.

C. Motion tracking

The tracking setup is comprised of a Raspberry Pi model 3B (RPI) and a Raspberry Pi Camera version 1.3. This bundle is capable of recording and transmitting images of 600×600 pixels at 90 frames per second. We connect the RPI via Ethernet to a Lenovo U430 Touch notebook, which we used to run the ROS core and perform image analysis.

We track the position of the sliding object on the plate by using a marker, as shown in Figure 2b. Because the marker is much darker than the test plates, we use simple binary thresholding to consistently track the position of the object.

A corner detection algorithm is used to determine the position of each corner of the tracking marker, which is a square.

Before starting the tilting process, a reference position of the test object is determined. During the tilting, the current observed position of the block is compared to the reference position, and if three out of four corners of the box surpass a threshold l then sliding is reported. The threshold l is defined as $l = 10 + \theta$ where θ is the current tilting angle of the platform. The angle θ is included to offset image drift caused by the gravitational forces pulling on the camera at larger angles.

IV. AUTOMATED ANISOTROPIC ACQUISITION

Acquisition experiments were performed using different combinations of materials for the test objects and the platform. The test objects include a pine wood cube with noticeable grain direction and a stainless steel nut; the platform surfaces are stainless steel and pine wood with smooth finish. Both test objects were used on both platform surfaces, resulting in four combinations: steel on steel, wood on wood, steel on wood, and wood on steel.

The materials used in our experiments, stainless steel and pine wood, have rough, dry and hard surfaces, making them suitable for our capture setup, and making the results of our work comparable to the reference literature. Stainless steel and other metals are expected to be mostly isotropic [16], whereas pine wood is typically considered to be anisotropic [17], [18]. We therefore expect elliptical friction cones from our dataset.

The ideal tilting motion would be in the direction that allows the object to slide the farthest as this most likely prevents touching a border. However, in our experiments we used only four tilting directions, each 90 degrees apart in order to produce reliable and smooth tilting paths. This could be revisited in future work with software updates or a specifically tailored tilting platform.

An algorithm was developed to act as controller to collect samples. First the arm is set to a leveled position. Then the tilting direction is selected based on the distance of the center of the tracked object from the borders of the recorded image. This is done in order to maximize the expected sliding distance. A weighted random choice is made which favors the longest distance. This allows the choice of a direction with shorter sliding distances, which avoids always tilting in the opposite direction of the previous slide, as this is typically the longest sliding distance.

Once a tilting direction is selected, we loop over an interval of tilting angles split into 50 steps. At first, we let this interval range from 0 to 67.5 degrees, which we note is a safe maximum tilting that we can use without worry of losing the test sample. We subsequently adjust the range to the mean plus minus three standard deviations from the mean after 25 samples have been gathered in a given direction. The reduced testing range enables smaller steps and therefore increases the data quality. Recall that a threshold l is used to detect sliding. At the moment of sliding, multiple data points are recorded over 300 milliseconds to create a noise reduced average. The

TABLE II: The most important features gathered for each sample. These fields define our protocol for recording data.

Name	Description	Unit
Euler angles	Oriental Euler angles of the end effector	Radian
Reference position	Pixel coordinates of each detected marker corner pre sliding	Pixel
Estimated variance	Estimated noise of the corner detection gathered during camera calibration	Pixel
Post sliding position	Pixel coordinates of each detected marker corner post sliding	Pixel
Tilting direction	The direction the platform was tilted in (Range 0-3)	Integer




Fig. 5: An illustration of the geometry and the relative vector orientations. Vector \mathbf{s} is in the plane spanned by \mathbf{d} and \mathbf{t} .

recorded features are listed in Table II. The full code of the tracking setup and the controller is open source [19].

To obtain the coefficient of friction μ , a common method is to use the angle of repose θ . This is the inclination angle of the surface at which the object starts sliding. From this, μ can be obtained by

$$\mu = \arctan \theta. \quad (1)$$

For anisotropic friction, the θ measured along the sliding direction \mathbf{s} is not necessarily the direction of maximum incline \mathbf{d} . The case $\mathbf{d} \neq \mathbf{s}$ occurs when there is a direction \mathbf{s} where the adjusted gravitational pull is higher than the resisting static friction. This can only happen with anisotropic friction.

We parameterize the tilting axis \mathbf{t} by β radians rotation of the world x unit axis \hat{i} around the world up z unit axis \hat{k} . This is illustrated in Figure 5. Let the axes be given by \hat{i} and \hat{k} , then we have

$$\mathbf{t} = \mathbf{R}(\hat{k}, \beta) \hat{i}. \quad (2)$$

The normal vector of the tilted plane can be obtained with

$$\mathbf{n} = \mathbf{R}(\mathbf{t}, \alpha) \hat{k} \quad (3)$$

where α is the inclination angle of the surface. The direction of the highest slope \mathbf{d} can be calculated as

$$\mathbf{d} = \mathbf{t} \times \mathbf{n}. \quad (4)$$

Observe that \mathbf{s} lies in the inclined plane. Therefore

$$\theta = \arcsin(s_z) \quad (5)$$

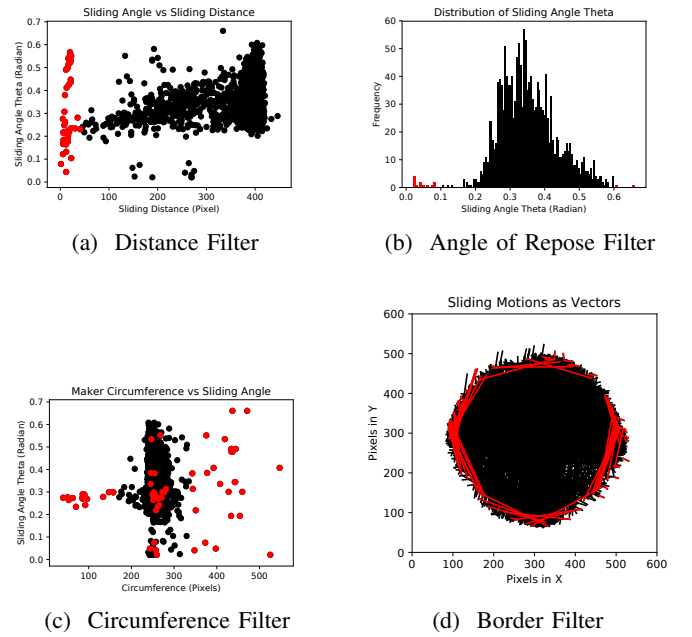


Fig. 6: These graphs visualize the filters we use for the experiment with a steel object on a steel surface. The red dots, bars or lines represent the samples that have been flagged by each filter.

where s_z is the z -component of the sliding direction. Because s_z is not measured due to the camera being mounted directly above the inclined plate, a 3D version of \mathbf{s} was obtained using

$$\mathbf{s} = \mathbf{R}(\mathbf{t}, \alpha) \mathbf{s}' \quad (6)$$

where \mathbf{s}' is \mathbf{s} as seen in the camera space.

We observed that the experiments did sometimes reveal some problematic behaviors. For example, objects sliding along the border could potentially result in a higher than average angle of repose because the favored sliding direction might be blocked. Additionally, the contact with the border materials might result in additional friction making these samples problematic. Therefore these samples were filtered based on pre- and post-sliding positions. Furthermore, a filter was employed based on the angle of repose where all samples which differed by more than three standard deviations from the average angle of repose were eliminated.

The marker based detection proved to be reliable but not perfect. In some cases, specular highlights on the tape and reflective metal caused some issues with the binary thresholding algorithm. To avoid false positives all sliding distances of less than 20 pixels on average were discarded. The circumference filter discarded all samples where either the pre or post sliding markers had a circumference that was more than three standard deviations from the average. A bigger or smaller circumference is a clear sign of a marker detection error.

All filters are illustrated on the example of the steel on steel surface pair in Figure 6.

TABLE III: The *Total* column contains the sum of the number of all filtered samples. The *Unique* column contains the number of unique filtered samples from each surface pair because some samples might be flagged from more than one filter.

	Unique	Total	Distance	Angle	Circumference	Border
Steel on Steel	138	181	50	14	35	82
Wood on Wood	49	55	21	6	4	24
Steel on Wood	228	233	84	27	20	102
Wood on Steel	137	143	7	10	5	121

We reconstruct the sliding direction relative to the objects orientation. The OpenCV library included with this particular ROS version did not include AR-Tag tracking. Therefore the reconstruction is based on the rectangular shape of the marker. A local coordinate system \hat{X} and \hat{Y} was aligned with the marker so that \hat{X} is parallel to the short side of the tracked rectangle. The local x-axis \hat{X} is always on the same line but can flip direction. The positive direction of \hat{X} was set to always point along the positive X -axis of the camera frame. With this assignment of the orientation of the axes, an adjustment had to be made to the gathered data: all relative sliding directions pointing in a negative \hat{Y} were mirrored, point wise. This allowed for symmetrical trends to be discovered and provided a view of at least 180 degrees.

V. EXPERIMENTS AND RESULTS

In total 6 330 samples were gathered, which took an average time of 22 seconds per sample. From these we removed 552 samples through filtering, which is 8.7% of the total data pool. More details on the filtered samples can be found in Table III. The full data set is available online [20].

The experiment includes four surface pairs of which three showed clear isotropic friction behavior, namely the steel object on steel plate, steel object on the wooden plate, and the wooden object on the steel plate. The isotropic behavior can be best seen in Figure 7. The figures for the respective surface pairs show noisy straight lines, which indicates that the friction does not differ significantly in a specific direction. While this noise could hide an anisotropic behavior, the interest lies in the exact frictional behavior of these materials. Therefore it is sufficient to assume that these surface pairs exhibit an isotropic friction behavior.

When comparing the average angle of repose between the intra-material surface pair with the inter-material pairs the difference is striking. The measured friction for the steel on steel surface pair is almost twice as high at 0.35 radian as any of the inter-material surface pairs, which are around 0.2 radian on average. This trend continues with the wooden object on the wooden plate which shows a higher friction. This could be due to the surface asperities having a more compatible pattern, which would result in a higher area of real contact and therefore a higher friction.

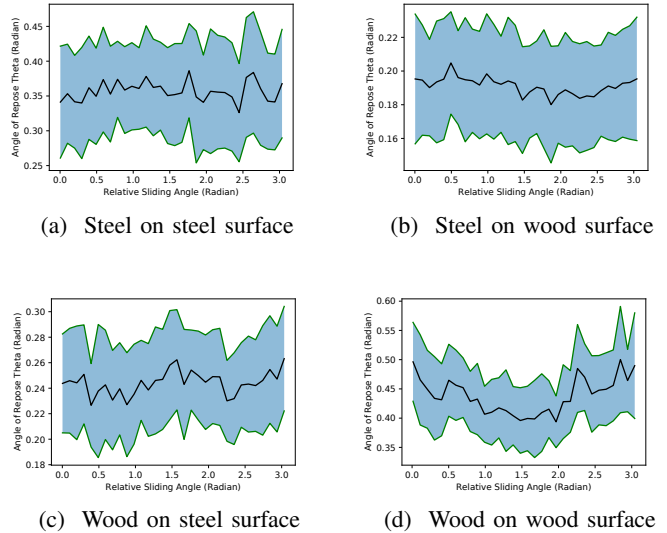


Fig. 7: Angles of repose θ for all directions relative to the object’s orientation. The black line is the average and the shaded area shows the standard deviation. The 180° range is divided into 32 bins with average and standard deviation computed for each bin. The wood on wood example stands out in particular as it shows clear anisotropic friction.

A closer look at the intra-material pairs reveals a difference of around 0.04 radians. This can be explained with the different treatments of the steel plate compared to the steel nut. The steel nut has been polished which leads to a smoother surface compared to the steel plate and therefore a slightly lower static friction is expected.

Another notable result is the variance of the steel on steel surface pair differs significantly. The shaded area in the graphs of Figure 7 display one standard deviation. For steel on steel it spans around 20% of the average angle of repose while all other surface pairs sit at around 10% of their average.

The heatmaps in Figure 8 show the angle of repose mapped to the pre-sliding position on the plate. While the lower friction, inter-material, surface pairs show a more equal distribution around the center of the plate, the intra-material ones did not. This indicates that the sliding on the inter-material surface pairs stopped before the object hit the opposing border. Additionally the low static friction observed in both of these could mean a potentially equally high or slightly higher dynamic friction which would explain the low sliding distances.

The only surface pair tested in this experiment to show anisotropic friction behavior is the wooden cube on the wooden plate. Anisotropic behavior was expected as the wooden cube has a tactile grain direction and the friction on the wooden plate is high enough to show these effects. The anisotropy is best visible in Figure 7d where the center has a lower angle of repose than both edges. Statistically, half of the results in this particular figure are mirrored, thus, the trend is most likely symmetrical. This follows our assumption

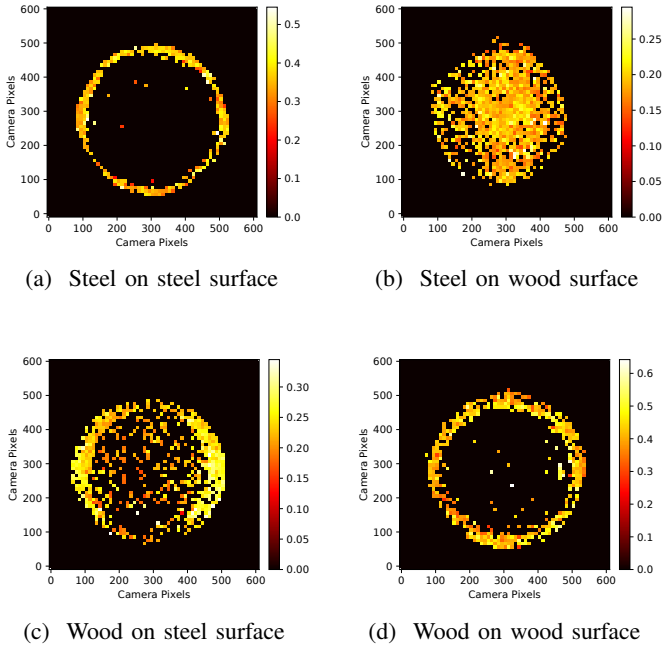


Fig. 8: The brighter a pixel the higher the value of θ at a given point. Black pixels represent no gathered data. A pixel on these heat maps is composed of a 10×10 pixel patch from the camera. Each subplot has its own color scaling. The most interesting result in this figure is the distribution of samples of the intra material surface pairs which is more equally spread than the case of inter material samples.

because the direction with lower friction is collinear to the grains direction.

It is useful to examine the results with the polar plots in Figure 9. Each dot is a single sample and each color represents a tilting direction. Because most of the surface pairs are isotropic, the distribution is equal among all directions. The wood on wood plot in Figure 9d shows a clear clustering of samples around the center, which is the low friction direction. This reveals a different set of problems. While this setup is able to obtain results for isotropic and some anisotropic cases, a strong anisotropic behavior would likely lead to a low sample count in high friction directions. This behavior is expected but poses a potential problem if a full 360 degree model is desired.

VI. DISCUSSION, LIMITATIONS AND CONCLUSION

We have proposed a system for automatically capturing anisotropic frictional responses that uses in-expensive components and a robotic arm with fully open source software for the experimental setup. Some preprocessing and filtering techniques have been applied which improve the quality of our database. However, we believe these could be circumvented by modifying the setup with better motion tracking capabilities. The extracted data is available in raw and preprocessed forms, and the presented analysis shows clear signs of anisotropic and

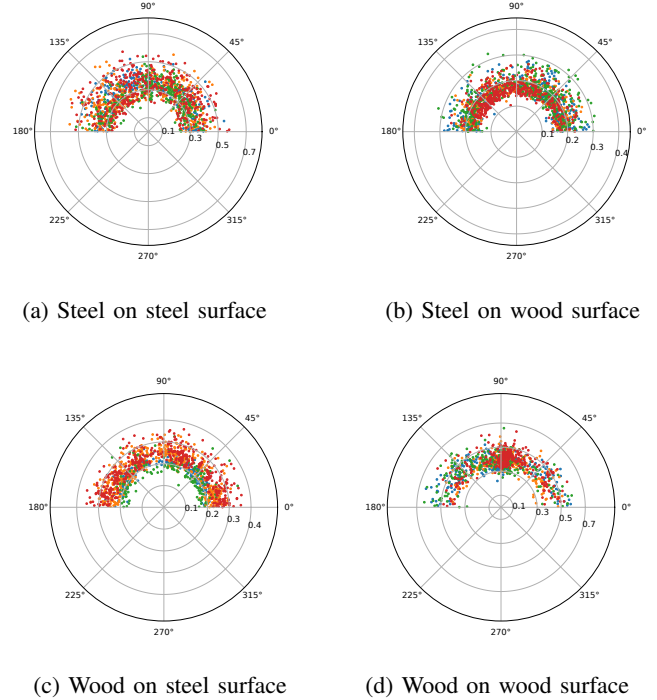


Fig. 9: The distance from the center displays the angle of repose θ in radian while the angle corresponds to the relative sliding direction of the object. Observe the wood on wood surface pair (d) clearly shows a clustering of samples in the low friction direction.

isotropic friction for combinations of two different materials. One contribution of our work is the development of a custom 3D end-effector that allows for the possibility of testing almost any desired combination of surfaces. Additionally, the setup is designed with reproducibility in mind, and we believe a wider robotics simulation community will be able to build on our experiences using our setup as their foundation.

Simulation friction models are often based on the concept of convex limit curves (e.g., an ellipse) and the principle of maximum dissipation [2], [21]. However, we observe that there are challenges to fitting our data to these type of models, with the main challenge being the ability to determine the local orientation of the limit curve as a parameter that depends on the local material properties. In our data set, the orientation appears to simply follow the test cube. However, no such guarantee can be claimed in general for all types of materials. Furthermore, fitting models for more complex surfaces may not be possible with our current experimental setup since there is no way to recover the relative orientation of the surfaces.

Anisotropy gives our measurements a real world bias as objects tend to slide in directions of least friction. This makes it hard to measure directions with large friction due to the unconstrained nature of our setup. However, if there is some overall shape or symmetry then the bias in sampling direction may still give sufficient information to create a complete cone.

Finally, our data reveals an interesting characteristic of

friction responses, which is that analytic functions are not realistic friction models. Rather, a fuzzy or stochastic model may be better able to predict real world frictional behavior. However, a challenge remains in finding a general approach to include this type of friction data in standard physics simulation pipelines.

ACKNOWLEDGEMENT

Thanks to Professor Inna Sharf, McGill university, for providing workshop, tools and robotic arm, and to Morten Engell-Nørregård, Evil Hippie Design, for the 3D model of the custom end-effector.

REFERENCES

- [1] N. Fazeli, R. Kolbert, R. Tedrake, and A. Rodriguez, "Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1437–1454, 2017.
- [2] Y. Jiang and C. K. Liu, "Data-augmented contact model for rigid body simulation," *arXiv preprint arXiv:1803.04019*, 2018.
- [3] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez, "More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 30–37.
- [4] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *European Conference on Artificial Life*. Springer, 1995, pp. 704–720.
- [5] N. Jakobi, "Running across the reality gap: Octopod locomotion evolved in a minimal simulation," in *European Workshop on Evolutionary Robotics*. Springer, 1998, pp. 39–58.
- [6] S. Carpin, M. Lewis, J. Wang, S. Balakirsky, and C. Scrapper, "Bridging the gap between simulation and reality in urban search and rescue," in *Robot Soccer World Cup*. Springer, 2006, pp. 1–12.
- [7] A. Zmitrowicz, "A theoretical model of anisotropic dry friction," *Wear*, vol. 73, no. 1, pp. 9–39, 1981.
- [8] S. Goyal, A. Ruina, and J. Papadopoulos, "Planar sliding with dry friction part I. limit surface and moment function," *Wear*, vol. 143, no. 2, pp. 307–330, 1991.
- [9] F. Marques, P. Flores, J. P. Claro, and H. M. Lankarani, "A survey and comparison of several friction force models for dynamic analysis of multibody mechanical systems," *Nonlinear Dynamics*, vol. 86, no. 3, pp. 1407–1443, 2016.
- [10] P. J. Blau, "The significance and use of the friction coefficient," *Tribology International*, vol. 34, no. 9, pp. 585–591, 2001.
- [11] M.-S. Kim, I.-Y. Kim, Y.-K. Park, and Y.-Z. Lee, "The friction measurement between finger skin and material surfaces," *Wear*, vol. 301, no. 1-2, pp. 338–342, 2013.
- [12] C.-G. Wallman and H. Åström, *Friction measurement methods and the correlation between road friction and traffic safety: A literature review*. Statens väg-och transportforskningsinstitut, 2001.
- [13] B. N. Persson, *Sliding friction: physical principles and applications*. Springer Science & Business Media, 2013.
- [14] R. D. Howe and M. R. Cutkosky, "Practical force-motion models for sliding manipulation," *The International Journal of Robotics Research*, vol. 15, no. 6, pp. 557–572, 1996.
- [15] D. K. Pai, K. v. d. Doel, D. L. James, J. Lang, J. E. Lloyd, J. L. Richmond, and S. H. Yau, "Scanning physical interaction behavior of 3d objects," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 2001, pp. 87–96.
- [16] E. Rabinowicz, "The determination of the compatibility of metals through static friction tests," *A S L E Transactions*, vol. 14, no. 3, pp. 198–205, 1971.
- [17] D. Atack and D. Tabor, "The friction of wood," *Proc. R. Soc. Lond. A*, vol. 246, no. 1247, pp. 539–555, 1958.
- [18] T. Ohtani, K. Kamasaki, and C. Tanaka, "On abrasive wear property during three-body abrasion of wood," *Wear*, vol. 255, no. 1-6, pp. 60–66, 2003.
- [19] K. Dressel, "Automated friction measurement package for ROS," <https://github.com/kenodressel/afm/>, September 2018.
- [20] —, "Anisotropic friction dataset," https://sid.erd.dk/wsgi-bin/lis.py?share_id=F1VXT6zzcd, September 2018.
- [21] J. Zhou, M. T. Mason, R. Paolini, and D. Bagnell, "A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation," *The International Journal of Robotics Research*, vol. 37, no. 2-3, pp. 249–265, 2018.