IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

All the Feels: A dexterous hand with large-area tactile sensing

Raunaq Bhirangi[1,2], Abigail DeFranco[*,1], Jacob Adkins[*,1], Carmel Majidi[1], Abhinav Gupta[1], Tess Hellebrekers[2] and Vikash Kumar[2]

Abstract—High cost and lack of reliability have precluded the widespread adoption of dexterous hands in robotics. Furthermore, the lack of a viable tactile sensor capable of sensing over the entire area of the hand impedes the rich, low-level feedback that would improve the learning of dexterous manipulation skills. This paper introduces an inexpensive, modular, and robust platform - the D'Manus - aimed at resolving these challenges while satisfying the large-scale data collection demands of deep robot learning paradigms. Studies on human manipulation point to the criticality of low-level tactile feedback in performing everyday dexterous tasks. The D'Manus comes with ReSkin sensing on the entire surface of the palm as well as the fingertips. We also demonstrate the generalizability of tactile models trained with the fully integrated system in a tactile-aware task - bin-picking and sorting. Code, documentation, design files, detailed assembly instructions, trained models, task videos, and all supplementary materials required to recreate the setup can be found on https://sites.google.com/view/dmanus

Index Terms—Multifingered Hands; Force and Tactile Sensing; Soft Sensors and Actuators

I. Introduction

HUMANS routinely operate in unstructured, cluttered environments through surprisingly imprecise, improvised motions. Think about finding the keys hiding at the bottom of your bag, pulling a box from the back of the fridge, or finding the steel ladle among the wooden spatulas. While you rely on vision to plan motion at a high level, executing low-level actions involves using a wealth of tactile signals to spatially understand and characterize the environment. The tactile information, combined with natural compliance and underlying motion, enables the effortless dexterity of the human hand. In moving towards robots with human-like sensorimotor abilities, there is a clear need for systems that integrate rich tactile sensing capabilities with dexterous motion.

However, the high dimensionality of dexterous systems also makes them difficult to control. Data-driven methods have emerged as promising approaches to high-dimensional control [1, 2], but success with dexterous manipulators has been limited [3, 4], and often restricted to simulation [5, 6]. The contact-rich nature of tasks like in-hand manipulation and tool use makes it difficult for policies learned in simulation to generalize to the real world. Collecting data from real-world interactions, on the other hand, is difficult due to the absence of an affordable, reliable hand that can handle the demands of large-scale data collection. Efforts aimed at developing such hardware have been few and far between [7, 8], largely due to the manufacturing cost and the lack of reliable sensing and actuation technologies.

Fig. 1: The D'Manus – a low-cost, 10 DoF, reliable prehensile hand with all-over ReSkin [9] sensing.

In this work, we leverage recent advancements in rapid prototyping, modular actuation and large-area sensing to present a hand that can make real-world dexterous learning accessible to a wider community of researchers and roboticists. Concretely, our contributions are as follows:

- We present the D'Manus – an inexpensive, robust prehensile hand geared towards real-world robot learning, complete with a detailed Mujoco-based simulation model for ease of development and prototyping. We rigorously test the hand to withstand long(>400)

Manuscript received: June 8, 2023; Revised: August 28, 2023; Accepted: October 12, 2023

This paper was recommended for publication by Editor Ashis Banerjee upon evaluation of the Associate Editor and Reviewers' comments.

[1]Raunaq Bhirangi, Abigail DeFranco, Jacob Adkins, Carmel Majidi and Abhinav Gupta are with Carnegie Mellon University, Pittsburgh, PA, USA

[2]Raunaq Bhirangi, Tess Hellebrekers and Vikash Kumar are with Meta

[*]equal contribution

[†]Correspondence: rbhirang@cs.cmu.edu,vikashplus@gmail.com

Digital Object Identifier (DOI): see top of this page.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

hours of operation with no breakages;
- We equip the D'Manus with customized, integrated ReSkin [10, 9] sensors that provide large-area tactile sensing over the entire surface of the palm and the fingertips, while maintaining sharp fingertips/nails critical for dexterous manipulation;
- We demonstrate the caliber of the D'Manus system along sensory effectiveness, dexterity, and robustness axes by learning tactile perceptive models for softness and texture identification; and validate their generalizability to unseen objects in a tactile-aware bin sorting task.

II. Related Work

A. Dexterous Hands and data-driven learning

The versatility of the human hand has long inspired a number of efforts aimed at creating similarly capable robotic hands dating back to the early days of robotics [11, 12]. Concurrent work in prosthetics and assistive robotics [13] has often overlapped with and contributed to research in creating general-purpose robotic hands. More recently, advances in material science and rapid prototyping as well as control algorithms have further pushed the envelope of capable dexterous hands [14]. Since these efforts have primarily been directed towards demonstrating added functionality on human control, they tend to fall short on the scalability, reliability, affordability, and other capabilities required for the prolonged operation demands of robot learning. Despite the recent advancements in data-driven robotics [1], robust dexterous platforms capable of meeting the data needs of real-world learning have been few and far in between [7, 8]. This has restricted recent investigations with dexterous hands to simulation [5, 6] or the few researchers who can afford the hardware expense [3]. The D'Manus is an open-sourced hand that fills a crucial void in the robot hand landscape – integrated large-area sensing and a palm unlike [7, 8], critical adduction and abduction capabilities unlike the Allegro, 30× less expensive than alternatives like [15, 16], and tested to be robust and easy to fix.

Additionally, most recent works aimed at solving dexterous manipulation [17, 3, 18] conspicuously use a single exteroceptive sensory modality – vision. Vision provides rich sensory information about the scene and the visual properties of objects, and has been successfully integrated with robot learning frameworks [1, 2]. However, dexterous tasks are generally contact-rich and require reasoning about contact information that cannot be captured entirely using vision. We posit that the lack of rich tactile information limits a manipulator's ability to effectively perform real-world dexterous manipulation tasks involving force control, flexible objects, and deformable media, particularly with smaller objects that receive degraded visual signals due to occlusions. The D'Manus comes with integrated large-area sensing that offers a rich tactile sensory modality and extensive spatial coverage suitable for learning such contact-rich manipulation skills.

B. Tactile sensing

The modality of touch has a long history in robotic grasping and manipulation [19]. Several different modalities like capacitive [20, 21], resistive [22], piezoelectric [23], magnetic [24, 10], audio [25] and MEMS [26]-based sensors have been explored as tactile sensing alternatives for robotics. With the recent success of deep learning, especially in computer vision, optical tactile sensors [27, 28] have emerged as the popular choice of tactile sensor, due to their high resolution as well as their compatibility with popular neural architectures (CNNs) for processing signals. Most of these solutions, however, have limitations that significantly impede their ability to serve as effective sensors for capable hands, which have strict requirements in terms of sensing, space, cost and robustness. Some of these sensors are bulky [27, 28] or need direct electrical connections between the circuitry and the interface [23, 22], resulting in design constraints that compromise on the manipulation abilities of the hand. Some others are either expensive [22] or difficult to fabricate [20, 21] and cannot be easily replaced, making them less suitable for large-scale data collection given the inevitable wear-and-tear that comes from frequent contact with a wide variety of objects. Yet other alternatives that are affordable and have suitable form factors for dexterous hands tend to lack the resolution, shear sensing [26, 29] required for fine-grained control. Manufacturing, cost, and reliability challenges only escalate with larger area sensing systems, such as the MIT Glove [30], hex-o-skin [29], and uSkin [16, 31] among others [20, 21, 26].

The class of works that come closest to our proposition are [16, 31] that use *uSkin* sensors to sensorize a dexterous hand and demonstrate application in object classification and manipulation tasks. However, *uSkin* uses macro-scale magnets embedded in elastomer as the sensing interface, which involves complex design to avoid crosstalk between magnetometers [24], is bulky, and creates an external magnetic field which can interfere with the environment. To counteract all of these problems, we turn to ReSkin [10, 9], which differs critically in the use of magnetic microparticles instead of macro-sized magnets. ReSkin offers the D'Manus a number of key advantages as a dexterous hand for robot learning, namely, (a) favorable form factor: ReSkin can be much thinner(∼2mm) than its closest alternatives (>5mm) enabling sharp fingernails critical to dexterous manipulation, (b) cost and replaceability: ReSkin is easily replaceable [9] and costs 50x lesser per sensor (∼ $20) than alternatives like uSkin (∼ $1000), and (c) wear resistance: the absence of a hard-soft interface within the elastomer significantly improves the durability of ReSkin [9] and, as a result, the D'Manus.

III. Platform and System Details

The D'Manus - a combination of Dynamixel and Manus, the Latin word for hand - is a low-cost, reliable prehensile robotic hand with immersive tactile sensing over its larger contact surfaces, i.e. the palm and fingertips as anatomized

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

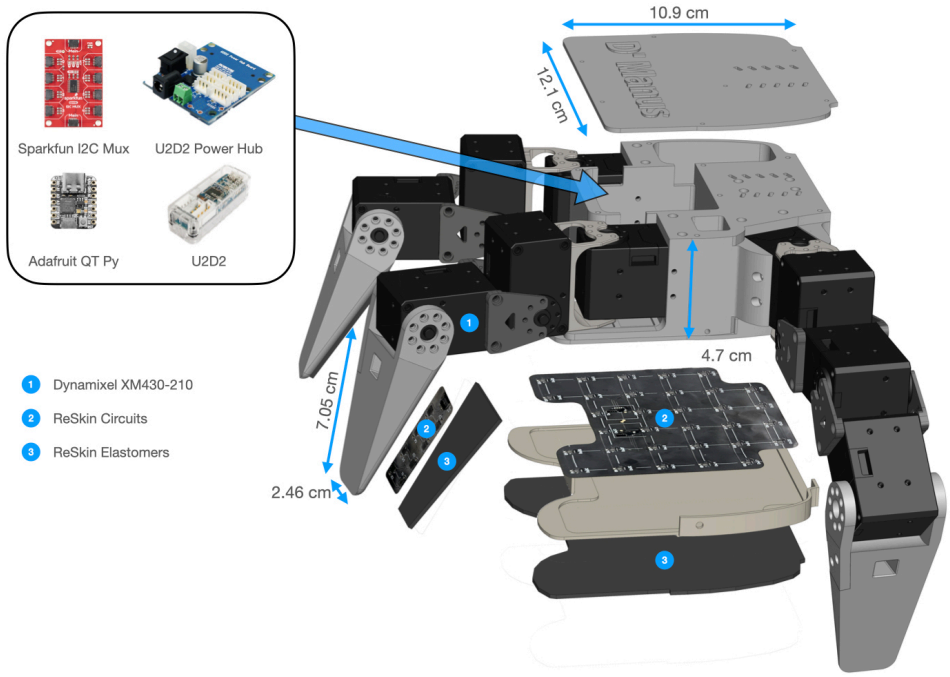

Fig. 2: Anatomy of the D'Manus hand: The D'Manus is actuated at joint level using Dynamixel XM430-210 smart actuators. ReSkin sensors are integrated with the fingertips and the palm. Each fingertip sensor is comprised of 8 magnetometers while the palm sensor consists of 32 magnetometers for a total of 56 magnetometers. Sensor and motor interfacing components are housed in the core of the hand.

in Fig. 2. To benefit the community and facilitate adoption, D'Manus is released as an open-sourced manipulation platform[1]. In this section, we detail the features and properties of the system.

Component	Cost
ReSkin Circuits	
Boards	$ 38.50
Assembly	$ 149.60
Parts	$ 56.00
Magnetic Microparticles	$ 5.50
Smooth-On DragonSkin-10 NV	$ 5.50
3D printed components	$ 50.00
Machined components	$ 200.00
ReSkin Interfacing	$ 20.45
Dynamixel Interfacing	$ 51.10
Dynamixel XM430-210 motors	$ 2899.00
Total	$ 3475.65

TABLE I: Cost breakdown for the D'Manus

A. The Hand: Construction and Interfacing

The D'Manus hand is a three-fingered, 10-DoF hand – each finger has three degrees of freedom, with a fourth DoF for the thumb. We select 3D printed parts and commercial actuators which allows the D'Manus to be easily customized and assembled while maintaining a low price point($3500), as detailed in Table I. The hand can be made compatible to be mounted on any robot arm or wrist attachment of choice using a simple 3D printed adaptor. In converging on the design of the D'Manus, we focus on three critical components: a palm, an opposable thumb and modular fingers, while ensuring that the fingertips can still perform a precision grasp. While we experimented with versions of the platform with up to 16 DoFs, we converged on the 10 DoF D'Manus as it strikes a balance between dexterity, cost, robustness, weight, and size.

[1]CAD models, bill of materials, circuit designs, assembly and setup instructions can be found on https://sites.google.com/view/dmanus

B. Large-area Exteroceptive Sensing: ReSkin

We use ReSkin [10, 9] to endow the hand with large-area exteroceptive tactile sensing. A ReSkin sensor is comprised of a magnetometer circuit in conjunction with a magnetic elastomer. Contact results in deformation of the elastomer which in turn results in a change in magnetic field that is picked up by the magnetometers. Drawing from [9], we scale the sensor circuits and the skins to the size of the palm and the fingertips while maintaining a thickness of 2mm for the skins. Each fingertip sensor is comprised of 8 magnetometers, while the palm sensor consists of 32 magnetometers. The signal from each magnetometer is the 3-axis magnetic flux density. Where our approach deviates significantly from [9] is an improved fabrication procedure for the magnetic elastomer skins used in this work. The skins are cured at room temperature without interfering magnetic fields, and then magnetized using a pulse magnetizer with a 4 Tesla (40 kOe) impulse. This change results in two improvements: (a) stronger signal strength: experiments with the circuits presented in [9] showed at least 5-6x stronger signal along each axis for the same deformation, and (b) ease and scalability of fabrication: magnetic grids scale poorly as the size of the skin increases, making them difficult to assemble

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

in grids as well as to pull apart post-curing Data from the sensors is streamed to the control computer via USB through a microcontroller + I2C mux. Fig. 2 illustrates the construction of the hand and how it integrates with the tactile sensors. A highlight of this design is also the large sensorized area of the palm ($\sim$ 11 cm x 12 cm) which facilitates stable power grasps and provides a base with force feedback for objects during in-hand manipulation tasks.

C. Control and Proprioceptive Sensing

To enable closed loop manipulation strategies with strong sensory feedback, the D'Manus also comes with a range of proprioceptive sensing capabilities at the actuated joints, as listed in Table II. Control strategies for manipulators lie on a spectrum between position/velocity control and force control. When interaction forces are negligible, position control enables more precise control of the end effector, while velocity control allows for smoother movements. On the other hand, constraints in the environment and frequent interaction forces lend themselves better to force control or "compliant" strategies [32]. The use of Dynamixel smart actuators afford the D'Manus a number of control modes as outlined in Table II, allowing operational flexibility for end user applications[2].

Property	Options
Control	Position, Velocity, Current, PWM
Proprioceptive Sensing	Position, Velocity, Current, Realtime tick, Trajectory, Input Voltage
Exteroceptive Sensing	ReSkin (30 Hz)
Limits	Position, Velocity, PWM, Current
Baudrate	9600 bps $\sim$ 4.5 Mbps

TABLE II: Operational Details for the D'Manus

D. Software

The D'Manus's software package includes a python driver that exposes all the operational modalities outlined in Table II, a detailed simulation model of the D'Manus based on MuJoCo (Figure 3), and a placeholder model of ReSkin sensors intended for prototyping. The software has been structured for ease of testing in simulation and transfer to the real hardware.

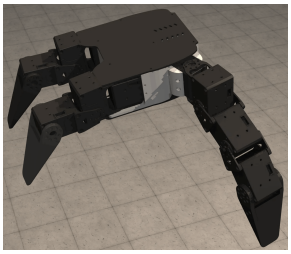

Fig. 3: Simulated D'Manus

IV. Experiments

The D'Manus is designed to sustain and support long hours of contact-rich interactions and data collection with minimal, easily fixable breakages. Such robustness allows the D'Manus to be used for long durations in a real-world robot learning setup similar to the systems demonstrated in ROBEL [7]. We evaluate the effectiveness of D'Manus as a testbed for real world robot learning along various axes –

1) Dexterity: In section VI-A, we evaluate D'Manus's prehensile ability by subjecting it to a variety of objects and grasping scenarios.
2) Tactile Perception: In sections VI-B, VI-C, we validate the discriminative ability of extensive tactile sensing by using the D'Manus to perform material, softness, and texture identification purely based on surface properties.
3) Perceptive Generalization: We demonstrate generalization of learned tactile models for softness and texture identification to unseen objects in section VI-C, and unseen tasks in section VI-D, to substantiate the stationarity and richness of ReSkin data.
4) Integrated system: We also corroborate the capabilities of the integrated D'Manus in section VI-D by exposing it to unseen, real-time interactions in a bin picking setup and demonstrating automated bin sorting purely from tactile information (no visual inputs).
5) Robustness: Finally, in section VI-E, we outline D'Manus's endurance and resilience towards extended periods of interaction rich operation.

While the D'Manus can be mounted on any robot arm, we used Franka Emika Panda robot for all our experiments. Neural network models presented in the following sections are are trained on a single GPU (NVIDIA GeForce GTX 1080 Ti), and never take more than 15 minutes of training time.

In the following section, we elaborate on the data collection and modeling choices for our learned tactile perception models before presenting experimental results in more detail.

V. Tactile Perception: Data and Modeling

We learn two generalizable tactile perception models training on ReSkin interaction from a variety of objects. The closest works in this space [16, 31] are restricted to classification problems that only demonstrate effectiveness on the same set of objects in the training set. We validate our models through testing on new, unseen objects. In this section, we detail the data collection setup and the modeling frameworks used to build these perception models.

A. Data Collection

To collect tactile interaction data, we fix the D'Manus such that the palm is facing upwards as shown in Fig. 4. For every object, we collect several trajectories of interaction data by placing it on the palm and executing a noisy, scripted motor babbling policy (at 30 Hz control frequency) for 10 seconds. While it is possible to train a

[2]PWM and current mode allow for force and hybrid position-force control

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

policy optimized for recognition, we selected motor babbling due to its simplicity and ability to collect data with minimal user intervention, and to study the effectiveness of the tactile sensors in isolation from policy learning. Our data collection trajectory was generated by smooth interpolation between randomly sampled waypoints in the joint space. The waypoints were sampled in permissible range respecting the joint limits and self-collision. We found that this naive policy provided sufficient data diversity for training our models, and demonstrating generalization to unseen objects and tasks.

Fig. 4: Data collection setup: Tactile data is collected by placing the object on the palm and executing a human-scripted interaction policy for motor babble.

As the interaction policy is executed, ReSkin data from the fingertips and the palm is streamed to the control computer at every time step. Each frame of data consists of 3-axis magnetic flux measurements for each of the 56 magnetometers enumerated in Fig. 2. Sample data from an interaction trajectory can be seen in Fig. 5. A more dynamic visualization of the raw data can be found in the accompanying video.

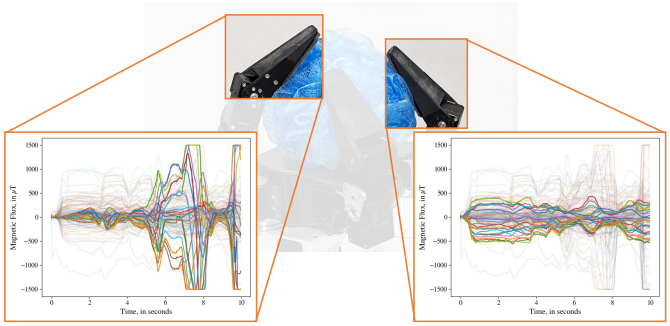

Fig. 5: Sample ReSkin data: Visualization of tactile data from two of the fingers while interacting with the loofah in Fig. 4.

B. Model Learning

Vision-based tactile sensors [27, 28] have naturally leveraged convolutional neural networks (CNNs) as a backbone for processing tactile information, as the signal is fundamentally images. In contrast, the electromagnetic signals of ReSkin have much less redundancy and have relatively lower dimensionality. In order to allow the learning algorithm to pick from a larger class of functions, we choose to use fully connected multilayer perceptrons(MLPs) as building blocks for the neural architecture used to process ReSkin signal. Further, contact information from interaction is naturally sequential, and our model architecture must be capable of leveraging temporal correlations. To ensure this, we use a recurrent neural architecture, an LSTM, at the base of our model. The neural architecture used in this work, as shown in Fig. 6, consists of an LSTM with 2 hidden layers followed by 3 fully connected layers. All the models in the experiments take a sequence of magnetic flux vectors $\mathbf{B}_{t\times 168}$ as input and are classification models trained to minimize cross-entropy loss $-\sum_{x\in X}\log P(f_\theta(x)=\hat{y})$.

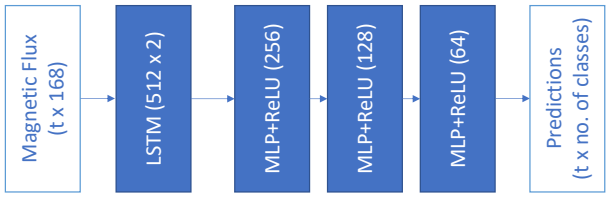

Fig. 6: Model architecture

VI. Results

A. Dexterity of the D'Manus

We qualitatively demonstrate the dexterous capabilities of the D'Manus (Fig. 8) using interactions with everyday objects. We observe that the D'Manus is effective at grasping and (in-hand as well as hand-arm) manipulation of day-to-day objects. Its abilities, however, are somewhat restricted for in-hand manipulation of small objects (e.g. counting coins on palm). This is in accordance with the dexterity and robustness trade-off we made and detailed in section III-A.

B. Tactile Perception: Material Identification

To establish the effectiveness of tactile perceptual capabilities, we task the D'Manus to leverage only its tactile signals to classify objects. The idea is to demonstrate that we can build tactile models capable of capturing the differences between the tactile signature of different materials as obtained by the D'Manus. We pick a set of six balls of identical shape and size, but with a different outer covering – small bubble wrap, large bubble wrap, corrugated cardboard, silicone sponge, a combination, and no covering material – as shown in Fig. 7, and collect interaction data as described in Sec. V-A. We use a 30-5 train-validation split for each ball.

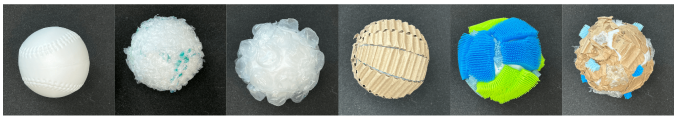

Fig. 7: Material coverings for Material Identification Task: Uncovered, small bubble wrap, large bubble wrap, corrugated cardboard, silicone sponge and combination of materials

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

Fig. 8: Illustration of the D'Manus grasping different objects using a variety of grasps [33, 34]

We train classification models to learn to predict a probability distribution over the six materials from tactile data, as detailed in section V-B. Our models show a 71.24% accuracy on the 30 held-out trajectories (5 per material) described above, as shown in Table III. This result confirms the discriminability of tactile interaction data obtained by the D'Manus.

Task	Validation accuracy
Material Identification	71.24%
Softness Classification	76.17%
Texture Classification	59.03%

TABLE III: The D'Manus can distinguish between different materials purely using tactile feedback (Sec. VI-B). Further, models trained for softness and texture classification generalize to interactions with unseen objects (Sec. VI-C).

C. Perceptive Generalization: Softness and Texture

Having verified the distinguishability of sensory signals of D'Manus, we shift our focus to the consistency of the sensory signal across different objects and scenarios. We develop tactile perception models for softness and texture identification and demonstrate its generalization to unseen objects. For generalization, it is imperative that the (a) D'Manus's sensors capture overall surface characteristics from interaction, (b) tactile perception models are robust to sensory drift over time, and (c) models are effective outside the training environment. To make our models robust to drift, training data is collected over an extended period of time (a few days). Test data is collected in the subsequent days to validate robustness to drift. Further, the bin sorting experiment in Sec. VI-D is performed about two weeks after training data is collected.

To learn tactile identification models, we would need quantifiable descriptions of surface characteristics. We create a three-point scale to quantify softness – Hard, Medium, Soft – as well as texture – Smooth, Medium, Rough. We manually assign softness and texture labels to over 50 objects by consensus among the authors before starting the study. We use a set of 20 training objects and 9 validation objects for these tasks. The full set of objects and their split can be found in Fig. 9. The corresponding datasets are created by collecting 15 trajectories of tactile interaction data for each of the training objects and 5 trajectories for each of the validation objects. We use the training data to train softness and texture identification models and examine their generalizability to unseen objects.

1) Softness Classification: We train a classification model as described in section V-B to predict softness categories, and present the results in Table III. Our models successfully learn to classify objects on the softness scale defined above.

2) Texture Classification: Training a model for texture classification analogous to softness classification has low generalization on validation set. This can be attributed to the difficulty of characterizing texture independently of its softness properties. A hard, highly textured ball (Fig. 9, bottom left) has better defined irregularities on the surface when compared to the ball of yarn (Fig. 9, bottom right). Without softness labels, neural networks struggle to find correlations between textured objects across softness categories.

To get around this problem, we train separate Softness-Conditioned texture identification models for each softness category trained on the corresponding subset of training data. To make a prediction, a sequence of tactile measurements is first input to the softness model which outputs a softness label. The texture model corresponding to this softness label is then used to predict a texture label from the same input sequence. The accuracy presented in Table III is the mean accuracy over the three softness categories.

A highlight of D'Manus design, as outlined in section III-B, is the large sensorized area of the palm. To corroborate this design choice, we train standalone softness and texture identification models corresponding to each individual fingertip and the palm. We evaluate the performance of these standalone models and present a comparison in Table IV. The performance on the individual finger models is significantly lower than the palm as well as the full prediction model. While some of this discrepancy can be attributed to the fingers losing contact with the object during parts of the interaction trajectory, strong performance of other models emphasizes the benefits of large area sensing available on D'Manus.

Component	Softness Accuracy	Texture Accuracy
Finger 1	56.27%	49.77%
Finger 2	48.48%	46.99%
Finger 3	59.59%	50.50%
Palm	74.31%	54.20%
All	76.17%	59.03%

TABLE IV: Comparison of models trained using data from different components of the hand.

To further substantiate the strength of tactile per-

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

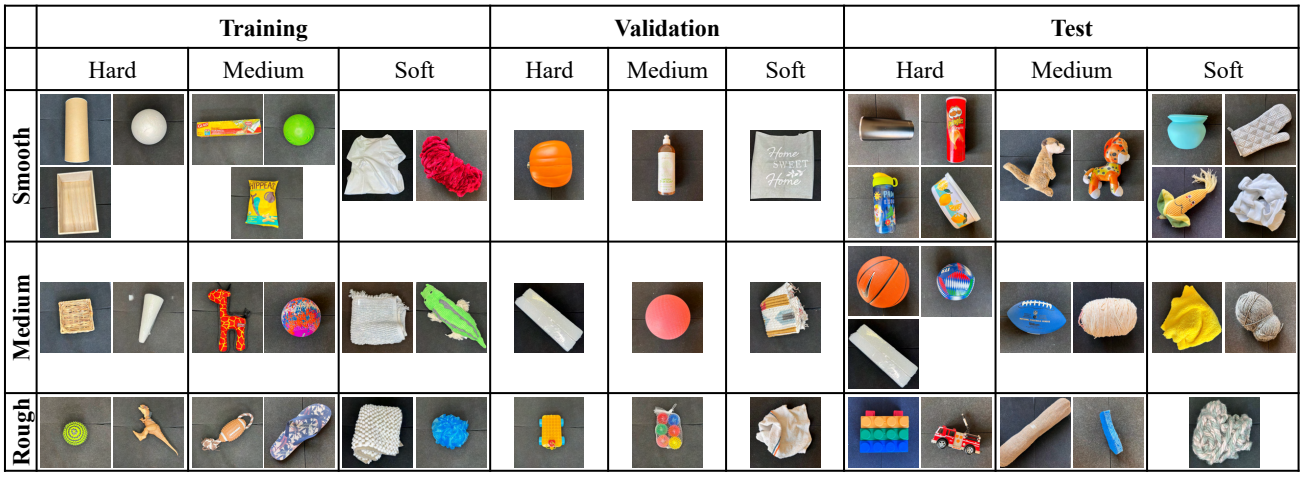

Fig. 9: Datasets used for Softness and Texture Identification Models

ceptual capabilities of the D'Manus, and demonstrate the working of the integrated system in an unseen environment, we deploy the learned softness and texture identification models in a tactile-aware bin sorting task in the following section.

D. Tactile Bin Sorting

As our final experiment, we access the ability of our trained models to generalize to realistic environments and tasks. For this evaluation, we pick a cluttered bin sorting experiment. We attempt to pick objects from a cluttered bin and sort them according to softness and texture from tactile signals. We start with a cluttered bin as shown in Fig. 1. The robot samples a random (x,y) location and reaches down into the bin until the ReSkin signal exceeds a specified threshold, indicating the presence of an object. Then, a predefined grasp is executed and the hand is raised. If the grasp is unsuccessful, the robot returns to random location selection. If the grasp is successful, we predict softness and texture labels and sort the object into corresponding bins. We then replace it by adding a new object to the bin and the process is continued. Over 20 successful grasps of different objects, our models achieve a prediction accuracy of 65% on both softness and texture prediction, confirming the ability of our models to extend to unseen tasks and environments in the real world. Further, it is worth noting that the tactile models trained with the hand upright are able to generalize to this setting where the hand is primarily operated in a downward-facing configuration. Through this experiment, we validate the ability of the integrated D'Manus system as well the generalizability of our models in performing tactile-rich tasks in unseen, real-world environments.

E. Robustness and Reliability

Amongst various versions of the platform, we have logged over 10,000 hours of operational time over the course of 12 months in 3 different locations with a total of 5 breakages. These breakages consisted of three motor failures, one 3D printed part failure, and operational deterioration of wires – all of which were repaired in-house within 30 minutes by computer scientists unfamiliar with the working details of the D'Manus. We attribute the robustness largely to the motor selection - the high-quality Dynamixel XM430-210 uses metallic gears and a high safety factor for our required force range. The version of the platform being released has significantly benefited from aggressive real world testing of prior versions. The specific copy of the D'Manus used for experimental results has been running for over 400 hours over the last 8 months with no breakages, corroborating our claims about the robustness and reliability of this system for real-world learning in contact-rich robotic tasks.

VII. Conclusions and Limitations

We present the D'Manus – a low-cost, 3D printable, prehensile robotic hand geared towards robot learning. The hand comes with multiple actuation modes, proprioceptive sensing abilities as well as ReSkin-based large-area tactile sensing. We demonstrate the dexterity of this platform in grasping a variety of objects. To exemplify the utility of the large-area sensing, we validate the discriminability of the tactile signal by learning models for material identification as well as category-level softness and texture identification. Further, we illustrate the transferability of learned tactile models to unstructured, real-world envirionments through a touch-based bin picking and sorting task. The design, assembly and setup instructions have all been open-sourced to facilitate adoption by the community.

Limitations: While we validate the tactile capabilities of the hand, it falls short of validating the tactile sensing in conjunction with dexterity. For future work, we would like to build a dexterous policy with tactile sensing to explore this further. The global semiconductor shortage limited the number of magnetometer chips we were able to integrate into this system. As a result, the system lacks sensing on the phalanges, ie. the surface of the motors, as well as the sides and backs of the fingertips. The design principles outlined in section III-B make this a simple extension of the present version and will be addressed in subsequent work. We also believe that unlocking the full

potential of all-over tactile sensing requires integration of other sensory modalities like vision and audio, allowing the system richer sensory inputs to solve complex dexterous tasks. Finally, this work lacks quantitative comparisons to other existing platforms due to the high cost involved in such a pursuit. We hope that open-sourcing D'Manus and its low cost will help with comparative evaluations with our system.

References

[1] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in 2016 IEEE International Conference on Robotics and Automation (ICRA), pp. 3406–3413, IEEE, 2016.

[2] C. Bodnar, A. Li, K. Hausman, P. Pastor, and M. Kalakrishnan, "Quantile qt-opt for risk-aware vision-based robotic grasping," in Proceedings of Robotics: Science and Systems, (Corvalis, Oregon, USA), July 2020.

[3] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, et al., "Learning dexterous in-hand manipulation," The International Journal of Robotics Research, vol. 39, no. 1, pp. 3–20, 2020.

[4] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk, K. Van Wyk, A. Zhurkevich, B. Sundaralingam, et al., "Dextreme: Transfer of agile in-hand manipulation from simulation to reality," arXiv preprint arXiv:2210.13702, 2022.

[5] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," arXiv preprint arXiv:1709.10087, 2017.

[6] T. Chen, J. Xu, and P. Agrawal, "A system for general in-hand object re-orientation," in Conference on Robot Learning, pp. 297–307, PMLR, 2022.

[7] M. Ahn, H. Zhu, K. Hartikainen, H. Ponte, A. Gupta, S. Levine, and V. Kumar, "Robel: Robotics benchmarks for learning with low-cost robots," in Conference on robot learning, pp. 1300–1313, PMLR, 2020.

[8] M. Wüthrich, F. Widmaier, F. Grimminger, J. Akpo, S. Joshi, V. Agrawal, B. Hammoud, M. Khadiv, M. Bogdanovic, V. Berenz, et al., "Trifinger: An open-source robot for learning dexterity," arXiv preprint arXiv:2008.03596, 2020.

[9] R. Bhirangi, T. Hellebrekers, C. Majidi, and A. Gupta, "Reskin: versatile, replaceable, lasting tactile skins," arXiv preprint arXiv:2111.00071, 2021.

[10] T. Hellebrekers, O. Kroemer, and C. Majidi, "Soft magnetic skin for continuous deformation sensing," Advanced Intelligent Systems, vol. 1, no. 4, p. 1900025, 2019.

[11] M. T. Mason and J. K. Salisbury Jr, "Robot hands and the mechanics of manipulation," 1985.

[12] G. A. Bekey, R. Tomovic, and I. Zeljkovic, "Control architecture for the belgrade/usc hand," in Dextrous robot hands, pp. 136–149, Springer, 1990.

[13] P. J. Kyberd and P. H. Chappell, "The southampton hand: an intelligent myoelectric prosthesis," Journal of rehabilitation Research and Development, vol. 31, no. 4, p. 326, 1994.

[14] C. Piazza, G. Grioli, M. Catalano, and A. Bicchi, "A century of robotic hands," Annual Review of Control, Robotics, and Autonomous Systems, vol. 2, pp. 1–32, 2019.

[15] A. Kochan, "Shadow delivers first hand," Industrial robot: an international journal, 2005.

[16] S. Funabashi, G. Yan, A. Geier, A. Schmitz, T. Ogata, and S. Sugano, "Morphology-specific convolutional neural networks for tactile object recognition with a multi-fingered hand," in 2019 International Conference on Robotics and Automation (ICRA), pp. 57–63, IEEE, 2019.

[17] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, "Deep dynamics models for learning dexterous manipulation," in Conference on Robot Learning, pp. 1101–1112, PMLR, 2020.

[18] H. Zhu, A. Gupta, A. Rajeswaran, S. Levine, and V. Kumar, "Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost," in 2019 International Conference on Robotics and Automation (ICRA), pp. 3651–3657, IEEE, 2019.

[19] B. Shih, D. Shah, J. Li, T. G. Thuruthel, Y.-L. Park, F. Iida, Z. Bao, R. Kramer-Bottiglio, and M. T. Tolley, "Electronic skins and machine learning for intelligent soft robots," Science Robotics, vol. 5, no. 41, p. eaaz9239, 2020.

[20] G. Cannata, M. Maggiali, G. Metta, and G. Sandini, "An embedded artificial skin for humanoid robots," in 2008 IEEE International conference on multisensor fusion and integration for intelligent systems, pp. 434–438, IEEE, 2008.

[21] T. Hoshi and H. Shinoda, "A large area robot skin based on cell-bridge system," in SENSORS, 2006 IEEE, pp. 827–830, IEEE, 2006.

[22] N. Wettels, V. J. Santos, R. S. Johansson, and G. E. Loeb, "Biomimetic tactile sensor array," Advanced Robotics, vol. 22, pp. 829–849, 2008.

[23] D. V. Dao, S. Sugiyama, S. Hirai, et al., "Analysis of sliding of a soft fingertip embedded with a novel micro force/moment sensor: Simulation, experiment, and application," in 2009 IEEE International Conference on Robotics and Automation, pp. 889–894, IEEE, 2009.

[24] T. P. Tomo, M. Regoli, A. Schmitz, L. Natale, H. Kristanto, S. Somlor, L. Jamone, G. Metta, and S. Sugano, "A new silicone structure for uskin—a soft, distributed, digital 3-axis skin sensor and its integration on the humanoid robot icub," IEEE Robotics and Automation Letters, vol. 3, no. 3, pp. 2584–2591, 2018.

[25] D. Gandhi, A. Gupta, and L. Pinto, "Swoosh!

rattle! thump!–actions that sound," arXiv preprint arXiv:2007.01851, 2020.

[26] L. U. Odhner, L. P. Jentoft, M. R. Claffee, N. Corson, Y. Tenzer, R. R. Ma, M. Buehler, R. Kohout, R. D. Howe, and A. M. Dollar, "A compliant, underactuated hand for robust manipulation," The International Journal of Robotics Research, vol. 33, no. 5, pp. 736–752, 2014.

[27] W. Yuan, S. Dong, and E. H. Adelson, "Gelsight: High-resolution robot tactile sensors for estimating geometry and force," Sensors, vol. 17, no. 12, p. 2762, 2017.

[28] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, et al., "Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation," IEEE Robotics and Automation Letters, vol. 5, no. 3, pp. 3838–3845, 2020.

[29] P. Mittendorfer, E. Yoshida, and G. Cheng, "Realizing whole-body tactile interactions with a self-organizing, multi-modal artificial skin on a humanoid robot," Advanced Robotics, vol. 29, no. 1, pp. 51–67, 2015.

[30] S. Sundaram, P. Kellnhofer, Y. Li, J.-Y. Zhu, A. Torralba, and W. Matusik, "Learning the signatures of the human grasp using a scalable tactile glove," Nature, vol. 569, no. 7758, pp. 698–702, 2019.

[31] S. Funabashi, T. Isobe, F. Hongyi, A. Hiramoto, A. Schmitz, S. Sugano, and T. Ogata, "Multi-fingered in-hand manipulation with various object properties using graph convolutional networks and distributed tactile sensors," IEEE Robotics and Automation Letters, vol. 7, no. 2, pp. 2102–2109, 2022.

[32] M. T. Mason, "Compliance and force control for computer controlled manipulators," IEEE Transactions on Systems, Man, and Cybernetics, vol. 11, no. 6, pp. 418–432, 1981.

[33] Y. Yang, C. Fermuller, Y. Li, and Y. Aloimonos, "Grasp type revisited: A modern perspective on a classical feature for vision," in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 400–408, 2015.

[34] T. Feix, J. Romero, H.-B. Schmiedmayer, A. M. Dollar, and D. Kragic, "The grasp taxonomy of human grasp types," IEEE Transactions on human-machine systems, vol. 46, no. 1, pp. 66–77, 2015.