

SpikeATac: A Multimodal Tactile Finger with Taxelized Dynamic Sensing for Dexterous Manipulation

Eric T. Chang^{*1}, Peter Ballentine^{*2}, Zhanpeng He^{*3}, Do-Gon Kim¹, Kai Jiang³, Hua-Hsuan Liang³, Joaquin Palacios¹, William Wang³, Pedro Piacenza¹, Ioannis Kymissis², Matei Ciocarlie¹




Fig. 1: **Overview of SpikeATac.** Left: completed finger and design. Right: integration on a parallel gripper and multifingered robot hand. Combining high-performance, taxelized PVDF for dynamic sensing with capacitive pads for static sensing, SpikeATac enables fast yet delicate manipulation. Used in conjunction with imitation learning and on-robot reinforcement learning fine-tuning, SpikeATac data can also enable multifingered dexterity such as in-hand reorientation even on fragile objects.

Abstract—In this work, we introduce SpikeATac, a multimodal tactile finger combining a taxelized and highly sensitive dynamic response (PVDF) with a static transduction method (capacitive) for multimodal touch sensing. Named for its ‘spiky’ response, SpikeATac’s 16-taxel PVDF film sampled at 4 kHz provides fast, sensitive dynamic signals to the very onset and breaking of contact. We characterize the sensitivity of the different modalities, and show that SpikeATac provides the ability to stop quickly and delicately when grasping fragile, deformable objects. Beyond parallel grasping, we show that SpikeATac can be used in a learning-based framework to achieve new capabilities on a dexterous multifingered robot hand. We use a learning recipe that combines reinforcement learning from human feedback with tactile-based rewards to fine-tune the behavior of a policy to modulate force. Our hardware platform and learning pipeline together enable a difficult dexterous and contact-rich task that has not previously been achieved: in-hand manipulation of fragile objects. Videos are available at roamlab.github.io/spikeatac.

I. INTRODUCTION

Tactile sensing typically falls into two categories: static sensing, for measuring sustained or slowly varying pressure, often with high *spatial resolution*, and dynamic sensing, for detecting rapid changes in pressure with high *temporal resolution*. These modes mirror biological mechanoreceptors,

with the human skin’s SA-I/II responding to sustained pressure and stretch and FA-I/II to dynamic changes.

These types of sensing are highly complementary. While static sensors can tell us about contact forces, contact geometry, and surface features, dynamic tactile sensors can reveal the precise moment when contact is made or broken, capture slip events, and measure vibrations created from extrinsic contacts and surface textures [1, 2]. Both are an important part of rich, multimodal tactile sensation.

Given the promise of dynamic sensing, the field has dedicated important resources to transduction methods that can convert this promise into advanced manipulation abilities. In particular, polyvinylidene fluoride (PVDF) has emerged as a particularly promising material for dynamic sensing. It is highly sensitive to dynamic events, offers high frequency bandwidth, is low cost, and can be manufactured in thin, flexible, and customizable form factors with custom taxel arrangements [3–6]. Furthermore, amplification circuits can even be fabricated directly on the PVDF film [6, 7], giving it a high ceiling for scaling to larger and integrated arrays.

In principle, these properties make PVDF well suited for creating dense, flexible, dynamic arrays. However, PVDF has almost exclusively been implemented in robot fingers as large, off-the-shelf taxels providing large-coverage, single-taxel sensing. Challenges remain in (a) how to build taxelized PVDF sheets that can integrate into a finger, (b) how to complement PVDF with a static modality, and (c) how to integrate its rich, difficult-to-simulate signals into modern,

^{*} indicates joint-first authorship. ¹ Dept. of Mechanical Engineering ² Dept. of Electrical Engineering ³ Dept. of Computer Science, Columbia University, New York, NY 10027, USA. This work was supported by a NASA Space Technology Graduate Research Opportunity and by a National Science Foundation Graduate Research Fellowship under Grant No. DGE-2036197. Correspondence email: eric.chang@columbia.edu

data-driven manipulation pipelines.

In this work, we address all three challenges with **SpikeATac**, a multimodal tactile finger containing a high frequency (4 kHz), 16-taxel PVDF array embedded at the contact surface, complemented by 7 capacitive pads embedded below the PVDF for pressure sensing. An optional accelerometer inside the finger provides additional high frequency feedback. At $45 \times 32 \times 25$ mm (length \times width \times thickness), SpikeATac is just larger than a human thumb, has a finger-shaped form factor conducive to stable contact states, and has 180° sensing coverage.

We demonstrate the utility of this new design in multiple ways. First, we show that SpikeATac enables rapid yet delicate manipulation, even of fragile objects, by leveraging the extreme responsiveness of PVDF to contact onset. Second, we integrate SpikeATacs into a four-finger robotic hand and show that raw SpikeATac signals can be incorporated into an imitation and reinforcement learning fine-tuning pipeline. We provide a recipe for learning on real, complex SpikeATac signals which achieves high dexterity even on fragile and delicate objects. Overall, our primary contributions are:

- To our knowledge, we are the first to propose a method for creating a robotic finger containing PVDF arrays with many individuated taxels across the entire surface. This design allows us to combine multi-taxel, dynamic, high frequency sensing with an additional transduction method (capacitive in our case) for static pressure.
- We leverage the extreme dynamic sensitivity of PVDF by complementing it with static sensing to achieve high-performance manipulation. We show an ability to perform very fast yet delicate manipulation of extremely fragile objects, an ability that, to our knowledge, has not been demonstrated before.
- We also show that highly sensitive, dynamic tactile sensing can be integrated into modern, learning-based motor dexterous policies. Specifically, we combine reinforcement learning from human feedback (RLHF) and tactile-based rewards to fine-tune the behaviors of a policy to precisely modulate force. We demonstrate this via a multimodal policy combining dynamic and static sensing to achieve dexterous, multifingered manipulation of fragile objects, an ability that, to our knowledge, has not been shown before.

II. RELATED WORK

1) *Dynamic Sensing*: Examples of dynamic modalities for measuring surface vibration include event-based piezoresistive arrays [8], capacitive sensors instrumented dynamically [2], hydrophones [9], MEMS microphones [10–14], and accelerometers or piezoelectrics embedded in the finger surface [15, 16]. We believe that high vibration signal dimensionality and spatial distribution are valuable, especially to a learned model or policy which has the ability to parse complex, overlapping signals and time-series data [8, 14, 17]. However, scaling these sensors to high-dimensional arrays is nontrivial due in part to transducer form factors, wiring, and sampling challenges.

2) *PVDF-based Fingertips*: PVDF is an established dynamic sensing modality in robotics, and previous work has shown its capability to sense contact events, slip events, and textures [2, 3, 15, 16, 18–20]. A number of works have combined PVDF strips with a static sensing modality in a fingertip [19, 21–23], however these works utilize only one or two PVDF strips, limiting their resolution. One study combined multi-taxel PVDF with a piezoresistive array [5], and another demonstrated how PVDF can be used for shear sensing when fabricated in dense multi-taxel modules [24]. Neither of these works, however, demonstrate multi-taxel PVDF in a fingertip.

3) *Multimodal Fingers*: Beyond PVDF, one can find many multimodal finger designs in the literature using a variety of sensing modalities. While many are flat form factors [25–28], we focus on multi-curved fingers that can sense diverse contacts during dexterous manipulation. Saloutos et al. built a spherical multimodal fingertip combining barometers and time-of-flight sensors [29], BioTac combines impedance sensing electrodes in both DC and AC configurations, temperature sensors, and a hydrophone for higher frequency vibration sensing [9], and TacTip and Minsound add a MEMS microphone to a camera-based tactile finger [10, 11]. Digit360 combines many modalities into a compact fingertip, including a camera, IMU, MEMS microphones, barometers, temperature sensors, and gas sensors [12].

All of these are valuable platforms for multimodal, dynamic manipulation. *The primary innovation of our work is the exploration of multi-taxel PVDF in a multimodal finger.* Multi-taxel PVDF provides distributed sensing across the finger, rather than only a few dynamic, high frequency (> 1 kHz) transducers. Lastly, we note that high bandwidth static sensors can also capture high frequency information. We favor a multimodal approach with designated dynamic sensors like PVDF so that the dynamic modality can be tuned to be sensitive to small transient events, while static sensors can preserve a larger dynamic range.

4) *On-robot RL fine-tuning*: RL has proven to be an effective approach for policy adaptation and improvement. A growing line of research [30–32] explores using RL to fine-tune pretrained policies directly on physical robots, thereby bridging the gap between offline training and real-world deployment. SERL [33], an open-source software suite, supports on-robot RL through classifier-based rewards and human-labeled sparse signals. DPPO [34], which fine-tunes diffusion-based policies on real hardware, demonstrates the potential of combining generative policy architectures with on-robot optimization. In this work, we integrate semi-sparse human reward labels with dense tactile-based rewards. This hybrid formulation leverages the efficiency of human guidance while grounding the learning process in rich contact information, enabling more reliable improvement of base policies during real-world fine-tuning.

III. SENSOR DESIGN

SpikeATac is a multimodal tactile fingertip. Its core feature is a 16-taxel PVDF array which provides very high sensitivity

to high frequency events, like making and breaking contact, extrinsic contact, and vibrations from surface textures. The PVDF is embedded near the surface of the finger's elastomer (Fig. 1) in order to be sensitive to these types of signals while being isolated from structural vibration. This design allows us to complement the dynamic PVDF by embedding a number of commercial off-the-shelf (COTS) capacitive taxels deeper under the surface in order to provide a static pressure response. Lastly, a 3-axis accelerometer acts as a secondary dynamic modality and is connected to a "nail" feature for surface exploration; however, we do not make use of this feature in this study.

PVDF produces charge proportional to applied strain. We use charge amplifiers to measure the charge generated at each electrode, therefore measuring applied strain. However, the feedback resistor of the charge amplifier provides a discharge path, creating high-pass filter behavior. A critical part of our work is that this configuration allows the PVDF to be very sensitive to transients without permanently saturating. See Section III-B for more electronics details.

The total material cost is US\$365, dominated by the COTS capacitive pads (US\$263 for seven such sensors alone).

A. PVDF Fabrication

The PVDF stack is a flexible, transparent piezoelectric film with metal on both sides forming 16 capacitor units. The taxelized side, fabricated using photolithography, contains pads, traces, and alignment marks. The common ground side, fabricated using a shadow mask, contains a single pad.

Each substrate is 4.5 cm x 4.5 cm x 100 μm thick pre-poled PVDF (PolyK). We sonicate the PVDF substrates for 15 min. in isopropyl alcohol and then rinse in deionized water to clean the surface before fabrication. A metal stack of 20 nm of a chromium adhesion layer and 200 nm of gold is deposited in a physical vapor deposition (PVD) system (Angstrom) with electron-beam evaporation. Then, we spin-coat positive photoresist (S1805, Dow) onto the sample. Using a photomask written with laser lithography (Heidelberg Systems), the photoresist is patterned by a mask aligner system (SÜSS). The photoresist pattern is developed in AZ 300 metal ion free developer, then transferred to the chrome/gold stack with two wet etches (Transene). This completes the taxelized side of the sample. We deposit the common ground plane in a PVD system through a shadowmask, with the same metal stack as the taxelized side.

To prepare the samples for SpikeATac fabrication, we cut the samples to the desired shape using alignment marks in the pattern. We use copper tape and a low-temperature curing silver paste (Micromax) to short the ground plane to a ground trace on the taxelized side which goes to the connector. Finally, the sensors are hot bar bonded to an flexible printed circuit (FPC) cable using an anisotropic conductive film (Hitachi) as the conductive interface. The FPC connects to a port on the printed circuit board (PCB) described in the next section.

B. Electronics Design

The "on-finger" electronics consist of a 2-PCB stack housed within the finger. The PCBs are connected with board-to-board

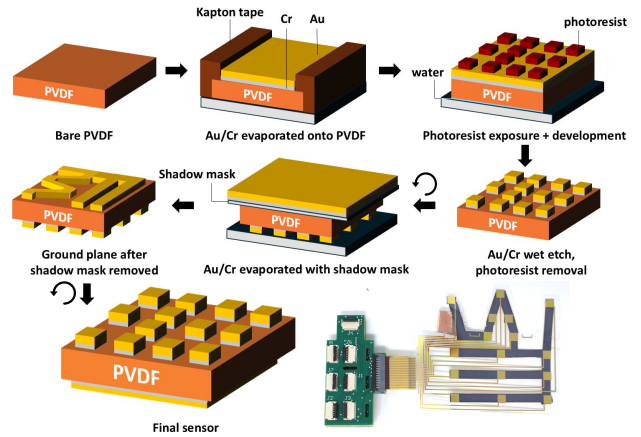


Fig. 2: The PVDF fabrication process along with a photo of the finished sensor attached to the PCB stack with an FPC cable.

connectors (Molex SlimStack) and handle all amplification and sampling. Flat flexible cable (FFC) ports interface with the capacitive sensor tails and the PVDF's FPC cable. A capacitance-to-digital converter (CDC, AD7147) samples the capacitive sensors. The stack also contains the three-axis accelerometer (MC3479).

PVDF's high impedance output requires careful electronics for sensitive, low-noise signals. Each of 16 PVDF signals are passed through a charge amplifier ($R_f = 1.2G\Omega$, $C_f = 22pF$) using a high impedance JFET op amp (AD8643), then sampled with an analog-to-digital converter (ADC, ADS7953), biased to 2.5 V. Our configuration provides a charge to voltage gain of $-\frac{1}{C_f} = -\frac{10^{12}}{22}$, high-pass filter cutoff frequency of $\frac{1}{2\pi R_f C_f} = 6.0$ Hz, and a time constant of $R_f C_f = 30$ ms.

The on-finger PCB stack is connected to an "off-finger" Teensy 4.1-based control board and communicates using serial peripheral interface (SPI). The control board controls up to 4 SpikeATacs and 4 Dynamixel chains and is what we use in all of our experiments. The Teensy communicates with a personal computer (PC) over USB using micro-ROS. When running only SpikeATacs, the PVDF is sampled at 4 kHz, the accelerometer at 1 kHz, and the capacitive sensors at 40 Hz.

C. Complete SpikeATac Fabrication

Finger fabrication contains 4 steps: mold preparation, skeleton preparation, casting, and PVDF layer application.

1) *Mold Preparation*: Half of a two-part mold (Clear V5 resin, Formlabs; washed and 15 min./60 $^{\circ}\text{C}$ UV cure) controls the front finger surface and contains a feature to create a 1.5 mm indentation on the finger surface for the PVDF. Once per mold, we use a brush to apply 2 layers of Inhibit X (Smooth-On) to prevent cure inhibition [35] (outgassing each coat for 30 min). The other half of the mold (PLA) controls the back of the finger geometry. At every use, we spray a thin coating of release agent (EaseRelease200) onto both molds.

2) *Skeleton Preparation*: A rigid skeleton forms the structure of the finger (Grey V5 resin, Formlabs; washed and 15 min./60 $^{\circ}\text{C}$ UV cure). After inserting pocket nuts, we adhere 7 capacitive sensors (Singletact, 8 mm/10 N model,

PPS) onto the skeleton using double-sided tape (300LSE, 3M). We route the Singtacts' 50 mm tails into the skeleton's internal space and connect them to the PCB's FFC ports. We apply 2 coats of Inhibit X to the skeleton (outgassing each coat for 30 min.), followed by 1 coat of silicone primer (SS4120, Momentive) to encourage bonding.

3) *Casting*: We attach a JST cable, carrying power and SPI buses to the control board, to the on-finger PCB. We suspend the skeleton in the mold by fastening the PLA negative mold to the back of the skeleton (M1.6 screws), which also fixes the PCB stack in its permanent position. We assemble the mold with screws and an O-ring. We mix and degas Ecoflex 00-50 elastomer (Smooth-On) and pour it into the mold, allowing it to cure at room temperature for 8 h.

4) *PVDF Layer Application*: Applying PVDF is nontrivial due to its imperfect fit on the finger's multi-curved surface and the difficulty of achieving strong PVDF/elastomer adhesion. We first adhere the PVDF to a flat, 0.5 mm thick Ecoflex 00-50 elastomer sheet of the same shape (Sil-Poxy, Smooth-On). After curing, we adhere this elastomer/PVDF stack into the finger surface's indentation (Sil-Poxy). We find that this 2 step process provides some forgiveness in the alignment to the finger surface and provides improved adhesion as the problematic elastomer/PVDF interface is cured completely while flat. Finally, we paint thin layers of Ecoflex 00-50 on top of the PVDF to form a thin layer (~3 coats).

IV. CHARACTERIZATION AND SENSITIVITY ANALYSIS

To characterize the signal response of the different modalities, we mount the sensor on an ATI Gamma 6-axis force/torque (F/T) load cell (SI-130-10 calibration) and probe it with a hemispherical indenter ($d = 6$ mm) on a linear probe at 1 mm/s and 10 mm/s. The finger is probed in 3 conditions: moderate touch (~3.5 N), light touch (probe just barely makes visual contact), and approach only (probe stops 0.5 mm away from finger). Sensor responses are shown in Fig. 3.

We note that, at the slow speed, PVDF provides a lower response and the signal dies away quickly. In other words, at this probe velocity, the high-pass filter behavior of the charge amps dominates the response. However, for faster incoming speed, the PVDF signal is rich in information: PVDF is very sensitive to the initial contact, spiking to saturation even at very low force. In fact, the PVDF signal starts to spike before the load cell or capacitive sensors move above their noise floor, showing that the PVDF's minimum detectable force at this probe velocity is below the load cell's noise floor (1 standard deviation of noise is 20 mN). The light touch condition also supports this conclusion, as the PVDF responds and the load cell does not.

The approach only data shows that the PVDF also exhibits a small response to object proximity, largely due to static charge, but this effect is small in comparison to true contact signals, even for light touch. We empirically observe that the proximity effect increases after adding the elastomer layer.

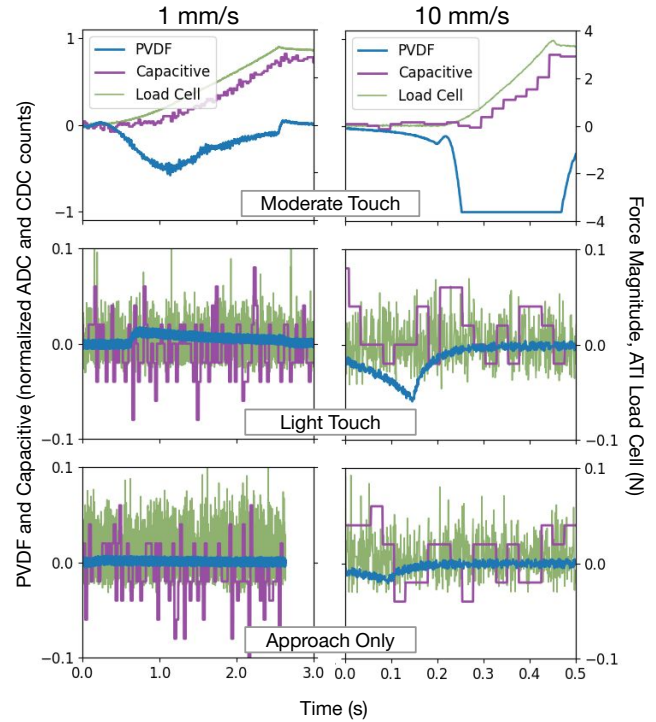


Fig. 3: PVDF, capacitive, and ATI Gamma responses when probing with a hemispherical indenter ($d = 6$ mm, see Fig. 4) at two linear probe speeds: 1 mm/s (left) and 10 mm/s (right). We test when under a moderate touch condition (~3.5 N), light touch condition (minimal contact established visually), and when approaching but not touching. Bottom-center PVDF and capacitive taxels are shown (location 8 on Fig. 4). ADC counts are normalized to the full measurement range; CDC counts are normalized to just above the moderate touch's maximum response. We note that, at the higher speed, PVDF can detect light contact that the other sensors cannot, or detect contact earlier than the other sensors.

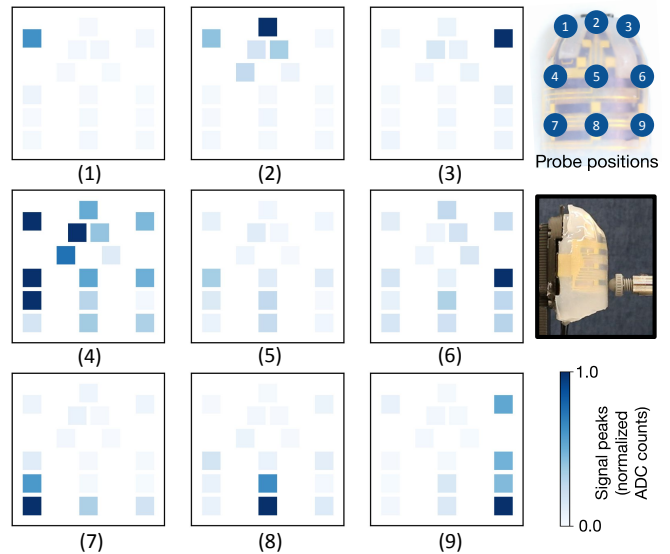


Fig. 4: Heatmaps of the maximum absolute value of the zeroed PVDF response from each taxel (i.e., a dark taxel represents an increasing or decreasing signal) depending on indentation location, showing the spatial information carried by the PVDF taxels.

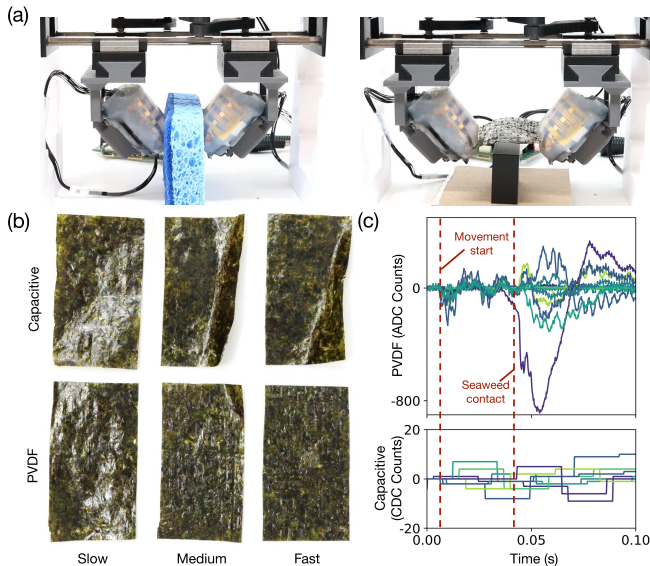


Fig. 5: (a) Setup of the fast and delicate grasping experiment, in which we command a *slow*, *medium*, or *fast* velocity to the gripper and stop after contact is detected, comparing PVDF-based and capacitive-based detection (Tab. I). (b) After 15 trials, the more fragile object (seaweed) is noticeably crumpled when using capacitive-based stopping, but appears nearly untouched in the PVDF condition at *fast* velocity. In fact, in the *medium* and *fast* conditions, the capacitive method often failed to detect contact at all. (c) Raw signals from one finger show a distinct PVDF spike at seaweed contact, while the capacitive signal does not reliably rise above its noise floor (each line/color shows an individual taxel from one of the two fingers).

The taxelized nature of our PVDF sensor is a key feature of this work. To test its behavior, we also probe the finger in 9 different locations to evaluate spatial information contained in the PVDF signals. These indentations are ~ 1.5 N and normal to the surface (via different finger mounts). A heatmap of the maximum absolute signals from each taxel (Fig. 4) shows that the PVDF provides spatial contact information. We note that the traces routed on the sample also contribute signal. For example, the indentation at position 9 (bottom right) activates the upper right taxel because the upper right taxel’s trace is routed near position 9 (see trace layout in Fig. 2). We note that position 5 has a weaker response due to a disconnected trace in the finger used for testing.

V. FAST AND DELICATE GRASPING

Utilizing PVDF’s ability to quickly and sensitively detect the onset of contact, we study SpikeATac’s ability to react fast enough to effectively stop on impact while moving at high speed. Moving fast while still being delicate is a valuable skill for robots moving quickly in unstable or safety critical environments and to grasp fragile objects without necessitating slow, time consuming movements. We thus compare how different sensing modalities enable this ability.

A. Experiment Setup

We mount two SpikeATacs on a custom parallel gripper containing a servo motor (XM430-210, Dynamixel) and a linkage mechanism controlling the linear motion (Fig. 5). We mount the gripper on a stand and place objects on a table between the fingers. During each grasp, we send a

velocity command to the Dynamixel’s velocity controller ($P=300$, $I=3500$), followed by a stop command (velocity = 0) once contact is detected (see below for details on contact detection). We check for the stop condition at 4 kHz (3 kHz for *slow* velocity due to Teensy memory limits), with PVDF and capacitive data updating at 4 kHz and 40 Hz, respectively. The PVDF has a median filter of width 3. The capacitive data is not filtered, however we remove occasional outlier values introduced by SPI bus noise.

We compare two methods of contact detection, one using PVDF, and the other using the capacitive sensors. In each case, we selected what we found to be the best performing method. For PVDF, we found a difference-based method to be the most effective, as the PVDF responds with sharp edges for the velocities investigated. However, we found this not to be the case for the capacitive signal, which does not respond as sharply, and we therefore selected a method which looks at the average signals compared to a static baseline.

- *PVDF Method*: Contact occurs once ≥ 2 PVDF taxels show a current–previous difference above a threshold.
- *Capacitive Method*: Contact occurs when the mean of the seven capacitive taxels moves beyond a threshold.

In each case, we chose thresholds empirically to be just above the noise floor and such that the motor doesn’t stop prematurely due to noise at the movement start.

The methods above are only for the *initial* contact detection of either finger. In both methods, to ensure sufficient grasp force for lifting the object, we follow the contact detection with a slow movement that stops when the mean of all capacitive values passes a second, object-specific threshold. For the PVDF and capacitive methods, respectively, we use a threshold of (40, 80) and (5.5, 6.5) counts for (finger 1, finger 2). We use a capacitive threshold of 3 (seaweed) and 5 (sponge) counts for the stable contact adjustment.

To run the experiment, we command one of three velocities to the motor and grasp one of two experimental objects: a sponge or seaweed (nori) sheets. Once contact is detected, the measured motor velocity at that timestep is recorded and the stop command is sent, followed by adjustment with capacitive feedback. We then record how far the gripper traveled after the ground truth contact point (inferred from object width), as well as the linear velocity of the gripper near the contact.

We run each of 12 experiments (two objects, two methods, three velocities) for 30 trials, verifying that the object is in a stable grasp at the end of each trial. Commanded velocities are the same for the two objects, however due to different object width, actual average achieved velocities were (91, 90) (*slow*), (180, 172) (*medium*), and (281, 180) mm/s (*fast*) for the (sponge, seaweed). We replace and remeasure the seaweed for each of six conditions.

B. Results

The sponge provides a quantitative comparison between the two modalities as it is deformable but firm enough that both methods can reliably grasp it gently (Tab. I). While there is little difference between methods at the *slow* velocity, the difference at higher velocity is apparent and consistent. PVDF

TABLE I: Distance traveled by the gripper (mean and stdev) beyond start of contact (smaller distance = faster stop) over 30 trials as illustrated in Fig. 5. We find that PVDF enables faster stopping, especially for the fragile object (seaweed), whereas the capacitive sensors are not able to reliably stop before crushing the object in the medium and fast conditions.

	Sponge			Seaweed		
Velocity	Slow	Medium	Fast	Slow	Medium	Fast
Cap. (mm)	0.8±0.4	1.9±0.7	3.7±0.5	13.3±2.3	–	–
PVDF (mm)	1.2±0.2	1.2±0.4	1.7±0.1	3.6±3.8	1.9±0.6	2.4±0.4

allows the gripper to respond faster upon initial contact at the high velocities than when using capacitive alone. These results are consistent with the expected behavior given that PVDF’s response is faster and more sensitive.

The results on the more fragile object (seaweed) are more striking. Over 30 trials at each speed, grasping with the PVDF sensors succeeded every time (0/30 crushed objects at each speed). In contrast, grasping with capacitive sensing succeeded at `slow` speeds (0/30 crushed objects) but failed often at `medium` and `fast` speeds (20/30 and 23/30 crushed objects, respectively). Additionally, the seaweed sheet after the high velocity PVDF condition appears nearly untouched. Due to the PVDF’s high pass filter behavior and detection algorithm, the faster we move with PVDF the more sensitive it will be. An illustration of the seaweed sheets after 15 trials is provided in Fig. 5.

We note that we cannot rule out proximity signals caused by static charge contributing to the PVDF’s contact detection. However, signals caused by proximity tend to be small relative to contact signals (Fig. 3), at least for the objects investigated in these experiments. The sharp edge and large magnitude in the raw PVDF data (Fig. 5) and consistent stopping behavior suggest the sensing of the contact event itself plays a critical role.

Finally, we use the fast grasping capability demonstrated here to grasp and pick up a wide set of delicate objects. See the [supplementary video](#) for examples.

VI. LEARNING DELICATE OBJECT IN-HAND ROTATION

So far, we have focused on standalone characterization of SpikeATac, as well as its integration with a parallel gripper for fast yet delicate manipulation. However, a traditional limitation of complex tactile sensors is the difficult integration in fully dexterous, multifingered manipulation. This is particularly true for current data-driven sensorimotor policies, for example for in-hand object re-orientation, which often train in simulation [36–38]: if a sensor provides highly complex signals which are difficult to simulate, how can it be used for learned motor control?

To address this challenge, we integrate SpikeATac into a dexterous, four-finger robot hand, and use a reinforcement learning from human feedback (RLHF) pipeline (Fig. 6) that is designed for fine-tuning behaviors on the real robot using raw sensor signals. To demonstrate SpikeATac’s sensitive and multimodal capabilities, we also tackle the problem of in-hand rotation via finger gaiting with only tactile and

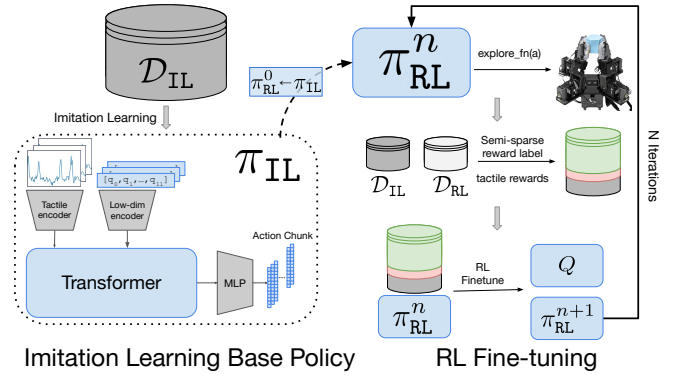


Fig. 6: Policy learning pipeline using on-robot RLHF in conjunction with imitation learning to fine-tune a control policy using real sensor data.

proprioceptive information. However, we make the problem even more difficult by extending it to fragile objects. To the best of our knowledge, this ability has not been previously shown in the literature.

A. On-Robot RL Fine-tuning for Behavior Improvement

At a high level, we start with a base policy π_{IL} that operates on raw sensor signals from our robot. To obtain π_{IL} , we train a behavioral policy in simulation via RL using a highly simplified, binary contact signal (touch vs. no-touch). We transfer this policy to the real robot using standard domain-randomization, and use it to collect a number of demonstrations on our real platform, which are in turn used to train a base policy π_{IL} via imitation learning.

While π_{IL} works well on hard objects (shown in our [supplementary video](#)), it immediately fails on fragile ones, as it has been trained with rigid objects and lacks the ability to produce delicate touch. We thus fine-tune it using RL, specifically Soft Actor-Critic [39], on raw sensor data acquired from our sensor. This pipeline shows how the highly-sensitive, difficult-to-simulate sensor signals can be integrated into learning-based manipulation.

To initialize RL fine-tuning, we set $\pi_{RL}^0 = \pi_{IL}$ and deploy it on the real robot to collect data. To enable effective data collection for training the Q-function, we introduce an exploration strategy by injecting Gaussian noise with small standard deviations into the distal joints. At each iteration, we use the most recent policy π_{RL}^n to collect a dataset \mathcal{D}_{RL}^n . For training, we aggregate this RL dataset with the original imitation learning dataset, i.e., $\mathcal{D}_{IL} \cup \mathcal{D}_{RL}^n$.

For RL fine-tuning, we incorporate two types of reward signals. The first type consists of task-specific rewards r_{task} provided by human annotators. For each trajectory, a human labels segments as “good” if the object is rotating and “bad” otherwise. We refer to this as semi-sparse reward labeling. Compared to sparse rewards, where only the final timestep is labeled based on whether the agent achieves the goal, this approach provides richer feedback across different parts of a trajectory. At the same time, unlike dense rewards, our reward labels do not require overly complex state estimation or engineering in the real-world setup.



Fig. 7: Rollout of the fine-tuned policy at three time points. The base policy $\pi_{\text{RL}}^0 = \pi_{\text{IL}}$ quickly crushes the paper object, while the fine-tuned policy using raw SpikeATac signals is capable of long rollouts without damaging the object.

The second type of reward signals is derived directly from tactile sensor observations. At a high level, our goal is to encourage the agent to reduce excessive contact forces (captured by capacitive readings) while increasing exploratory contacts (reflected in a higher number of PVDF spikes, i.e., making and breaking contact with the object). Formally, the tactile reward is defined as:

$$r_{\text{tac}} = w_1 \sum_i^{i \leq 4} \mathbb{1}[o_{\text{cap}}^i \geq 0.9] + w_2 \sum_i^{i \leq 4} \mathbb{1}[|o_{\text{pvdf}}^i| > 0.8]$$

Here, w_1 and w_2 are constants ($w_1 < 0$, $w_2 > 0$). o_{cap} and o_{pvdf} are tactile sensor readings that are normalized to the range of $[0, 1]$ using a sigmoid function, then scaled and shifted to the range $[-1, 1]$. i is the finger index. The combined reward for RL finetuning is $r = r_{\text{tac}} + r_{\text{task}}$.

B. Observation and Action Space

We carefully handle raw PVDF sensor signals in the observation because we do not want to discard high frequency information in the sensor data. We sample the PVDF at 450 Hz and capacitive sensors at 250 Hz (capacitive internal update rate is still 40 Hz), while the policy runs at 20 Hz (sampling frequencies are limited by the bandwidth of the microcontroller controlling the entire hand). The signals have a median filter of width 5. A history buffer of 64 PVDF and capacitive readings incorporates higher frequency information into the policy at each lower frequency timestep.

The observation contains the length 64 history buffer of sensor signals (16 PVDF and 7 capacitive signals per finger) and the 12 joint encoder readings. The action space is the hand’s 12 joint position set points for its 12 degrees of freedom (1 roll joint and 2 flexion/extension joints per finger).

C. Results

To evaluate performance, we roll out the policy at each iteration on a paper hexagonal prism ($d = h = 40$ mm) and paper cylinder ($d = 40$ mm, $h = 60$ mm) for 5 rollouts each. We evaluate amount of rotation and destruction rate for each policy (Fig. 8).

The initial $\pi_{\text{RL}}^0 = \pi_{\text{IL}}$ policy almost immediately crushes the objects (Fig. 7), however through this learning pipeline the policy learns to make visibly softer contacts over the learning iterations, as evidenced by the object state after the experiments (Fig. 8), while amount of rotation also increases.




Fig. 8: Mean and median rotation amount and object destruction rate (# times crushed out of 10) with each RL fine-tuning learning iteration evaluated on a paper hexagonal prism and cylinder (figure shows combined results). The images show examples of objects after manipulation at each iteration. The shaded region shows the interquartile range (IQR).

See the [supplementary video](#) for visual results of improved rotation over learning iterations.

We compare RL fine-tuning results to an IL baseline that is trained on all available demonstration data after all fine-tuning iterations, with the IL baseline achieving 5.7/2.9 (mean/mdn.) radians rotation and 0.7 destruction rate, notably lower than the RL fine-tuning results. We emphasize that the data we use for the IL fine-tuning baseline is usually hard to obtain in an IL setting, given the quick failures of the base policy. Here, we filter the RL fine-tuning dataset using human labels and only use the good trajectories for imitation learning.

These experiments show that this system and pipeline as a whole enable this difficult, delicate task. We note that the resulting policy in these experiments may be overfit to these particular objects and object sizes, and that we also don’t have a definitive analysis of what each sensor contributes. Our primary takeaway, however, is that this pipeline, which is designed for a real robot using real sensor signals in the reward, allows the use of SpikeATac data for new data-driven capabilities, despite the sensor being, as of now, very difficult to simulate.

VII. DISCUSSION AND CONCLUSIONS

Relative to other tactile sensing transducers, PVDF is a less prevalent choice of tactile modality in the literature despite its impressive sensitivity. This may in part be due to its susceptibility to noise, as a very high impedance system, and its cumbersome fabrication processes—the clean room tools used in this work are not readily accessible to many robotics-focused researchers.

A differential amplifier design and more advanced shielding will likely improve the PVDF’s performance [40]. To improve connector robustness, we’d also like to eliminate the ACF by connecting PVDF electrodes directly to a flexPCB [40]. We note that in addition to sensitivity to static charge, our sensor is also sensitive to how the system is grounded, but we found that with proper grounding the sensor signals were robust enough to enable all the experiments presented here.

Additionally, although the clean room processes used in this work make “democratizing” this sensor for the research community challenging, they are standard fabrication processes and should not be feared at scale. Screen printing may also be another avenue for scalable fabrication. In this vein, with careful implementation, we see PVDF as a compelling tool for many-taxel dynamic tactile sensing.

Overall, we believe SpikeATac is a powerful platform for studying multimodal tactile sensing and manipulation. The sensor’s many-taxel PVDF enables fast and delicate grasping. In conjunction with on-robot RL fine-tuning, we can integrate its raw, difficult-to-simulate multimodal sensor signals into complex learning-based policies that show unprecedented capabilities, such as in-hand reorientation of fragile objects. We believe that the combination of static and dynamic tactile sensing, integrated into capable hands and used in conjunction with powerful sensorimotor learning methods, shows great promise for future advances in robot dexterity.

VIII. ACKNOWLEDGEMENTS

We thank Trey Smith, Brian Coltin, Amr El-Azizi, and Rajinder Singh Deol for insightful discussion.

REFERENCES

- [1] R. S. Johansson and J. R. Flanagan, “Coding and use of tactile signals from the fingertips in object manipulation tasks,” *Nat. Rev. Neurosci.*, vol. 10, no. 5, pp. 345–359, May 2009.
- [2] M. R. Cutkosky and J. Ulmen, “Dynamic Tactile Sensing,” in *The Human Hand as an Inspiration for Robot Hand Development*, R. Balasubramanian and V. J. Santos, Eds. Cham: Springer International Publishing, 2014, pp. 389–403.
- [3] Z. Kappassov, J.-A. Corrales, and V. Perdereau, “Tactile sensing in dexterous robot hands — Review,” *Robot. and Auton. Systems*, vol. 74, pp. 195–220, Dec. 2015.
- [4] J.-H. Bae and S.-H. Chang, “PVDF-based ferroelectric polymers and dielectric elastomers for sensor and actuator applications: a review,” *Funct. Compos. Struct.*, vol. 1, no. 1, p. 012003, Mar. 2019.
- [5] C. Yu, M. R. Cavallari, and I. Kyriassis, “Tri-modal thin-film flexible electronic skin to augment robotic grasping,” in *2018 IEEE Micro Electro Mechanical Systems*, Jan. 2018, pp. 886–888.
- [6] Y.-J. Hsu, Z. Jia, and I. Kyriassis, “A Locally Amplified Strain Sensor Based on a Piezoelectric Polymer and Organic Field-Effect Transistors,” *IEEE Trans. on Electron Devices*, vol. 58, no. 3, Mar. 2011.
- [7] R. Dahiya, *et al.*, “Design and fabrication of posfet devices for tactile sensing,” in *TRANSDUCERS 2009 - 2009 International Solid-State Sensors, Actuators and Microsystems Conference*, June 2009, pp. 1881–1884.
- [8] T. Taunayazov, *et al.*, “Extended Tactile Perception: Vibration Sensing through Tools and Grasped Objects,” in *2021 IEEE Int. Conf. Intell. Robots Syst.*, Sept. 2021, pp. 1755–1762.
- [9] J. A. Fishel and G. E. Loeb, “Bayesian Exploration for Intelligent Identification of Textures,” *Front. Neurobot.*, vol. 6, June 2012.
- [10] N. Pestell and N. F. Lepora, “Artificial SA-I, RA-I and RA-II/vibrotactile afferents for tactile sensing of texture,” *Journal of The Royal Society Interface*, vol. 19, no. 189, p. 20210603, Apr. 2022.
- [11] I. Andrussow, *et al.*, “Adding internal audio sensing to internal vision enables human-like in-hand fabric recognition with soft robotic fingertips,” in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, Seoul, Korea, Sept. 2025.
- [12] M. Lambeta, *et al.*, “Digitizing Touch with an Artificial Multimodal Fingertip,” Nov. 2024, arXiv:2411.02479 [cs].
- [13] E. T. Chang, *et al.*, “An Investigation of Multi-feature Extraction and Super-resolution with Fast Microphone Arrays,” in *2024 IEEE Int. Conf. on Robot. and Autom.*, May 2024, pp. 3388–3394.
- [14] M. J. Yang, *et al.*, “Touch Classification on Robotic Skin using Multimodal Tactile Sensing Modules,” in *IEEE Int. Conf. Robot. Autom.*, May 2023, pp. 9917–9923.
- [15] R. Howe and M. Cutkosky, “Dynamic tactile sensing: perception of fine surface features with stress rate sensing,” *IEEE Trans. on Robot. and Autom.*, vol. 9, no. 2, pp. 140–151, Apr. 1993.
- [16] J. Son, E. Monteverde, and R. Howe, “A tactile sensor for localizing transient events in manipulation,” in *1994 IEEE Int. Conf. on Robot. and Autom.*, May 1994, pp. 471–476 vol.1.
- [17] P. Piacenza, *et al.*, “A Sensorized Multicurved Robot Finger With Data-Driven Touch Sensing via Overlapping Light Signals,” *IEEE/ASME Trans. Mechatron.*, vol. 25, no. 5, pp. 2416–2427, Oct. 2020.
- [18] R. Howe, *et al.*, “Grasping, manipulation, and control with tactile sensing,” in *IEEE Int. Conf. Robot. Autom.*, May 1990, pp. 1258–1263 vol.2.
- [19] X. Shi, Y. Wang, and L. Qin, “Surface Recognition With a Bioinspired Tactile Fingertip,” *IEEE Sensors J.*, vol. 23, no. 16, pp. 18 842–18 855, Aug. 2023.
- [20] Q. Guo, G. A. Al, and U. Martinez-Hernandez, “VibroTact: Soft Piezo Vibration Fingertip Sensor for Recognition of Texture Roughness via Robotic Sliding Exploratory Procedures,” *IEEE Sensors Letters*, vol. 8, no. 7, pp. 1–4, July 2024.
- [21] B. Choi, S. Lee, H. R. Choi, and S. Kang, “Development of Anthropomorphic Robot Hand with Tactile Sensor : SKKU Hand II,” in *IEEE Int. Conf. Intell. Robots Syst.*, Oct. 2006, pp. 3779–3784.
- [22] D. Goger, N. Gorges, and H. Worn, “Tactile sensing for an anthropomorphic robotic hand: Hardware and signal processing,” in *2009 IEEE Int. Conf. Robot. Autom.*, May 2009, pp. 895–901.
- [23] A. Ke, *et al.*, “Fingertip Tactile Sensor With Single Sensing Element Based on FSR and PVDF,” *IEEE Sensors J.*, vol. 19, no. 23, Dec. 2019.
- [24] P. Yu, *et al.*, “Flexible Piezoelectric Tactile Sensor Array for Dynamic Three-Axis Force Measurement,” *Sensors*, vol. 16, no. 6, p. 819, June 2016.
- [25] Z. Lu, X. Gao, and H. Yu, “GTac: A Biomimetic Tactile Sensor with Skin-like Heterogeneous Force Feedback for Robots,” *IEEE Sensors J.*, vol. 22, no. 14, pp. 14 491–14 500, July 2022, arXiv:2201.12005 [cs].
- [26] B. M. R. Lima, *et al.*, “Dynamic Tactile Exploration for Texture Classification using a Miniaturized Multi-modal Tactile Sensor and Machine Learning,” in *2020 IEEE Int. Syst. Conf.*, Aug. 2020, pp. 1–7.
- [27] U. Martinez-Hernandez and T. Assaf, “Soft Tactile Sensor With Multimodal Data Processing for Texture Recognition,” *IEEE Sensors Letters*, vol. 7, no. 8, pp. 1–4, Aug. 2023.
- [28] J. M. Romano, *et al.*, “Human-Inspired Robotic Grasp Control With Tactile Sensing,” *IEEE Trans. Robot.*, vol. 27, no. 6, pp. 1067–1079, Dec. 2011.
- [29] A. SaLoutos, *et al.*, “Design of a Multimodal Fingertip Sensor for Dynamic Manipulation,” in *2023 IEEE Int. Conf. Robot. Autom.*, May 2023, pp. 8017–8024.
- [30] L. Smith, Y. Cao, and S. Levine, “Grow your limits: Continuous improvement with real-world rl for robotic locomotion,” in *2024 IEEE Int. Conf. Robot. Autom.*
- [31] L. Smith, I. Kostrikov, and S. Levine, “A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning,” *arXiv preprint arXiv:2208.07860*, 2022.
- [32] Z. Zhou, *et al.*, “Efficient online reinforcement learning fine-tuning need not retain offline data,” *arXiv preprint arXiv:2412.07762*, 2024.
- [33] J. Luo, *et al.*, “Serl: A software suite for sample-efficient robotic reinforcement learning,” 2024.
- [34] A. Z. Ren, *et al.*, “Diffusion policy policy optimization,” in *arXiv preprint arXiv:2409.00588*, 2024.
- [35] B. Venzac, *et al.*, “PDMS Curing Inhibition on 3D-Printed Molds: Why? Also, How to Avoid It?” *Anal. Chem.*, vol. 93, no. 19, pp. 7180–7187, May 2021.
- [36] G. Khandate, *et al.*, “R \times R: Rapid eXploration for Reinforcement learning via sampling-based reset distributions and imitation pre-training,” *Auton Robot*, vol. 48, no. 7, p. 17, Aug. 2024.
- [37] J. Wang, *et al.*, “Lessons from Learning to Spin ‘Pens,’” Oct. 2024, arXiv:2407.18902 [cs].
- [38] H. Qi, *et al.*, “General In-Hand Object Rotation with Vision and Touch,” Sept. 2023, arXiv:2309.09979 [cs].
- [39] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *Int. Conf. on Mach. Learn.*, 2018.
- [40] A. J. Yeiser, *et al.*, “The UmboMic: a PVDF cantilever microphone*,” *J. Micromech. Microeng.*, vol. 34, no. 8, p. 085010, July 2024.