

Active Exploration for Real-Time Haptic Training

Jake Ketchum¹, Ahalya Prabhakar², Todd D. Murphey¹

Abstract—Tactile perception is important for robotic systems that interact with the world through touch. Touch is an active sense in which tactile measurements depend on the contact properties of an interaction—e.g., velocity, force, acceleration—as well as properties of the sensor and object under test. These dependencies make training tactile perceptual models challenging. Additionally, the effects of limited sensor life and the near-field nature of tactile sensors preclude the practical collection of exhaustive data sets even for fairly simple objects. Active learning provides a mechanism for focusing on only the most informative aspects of an object during data collection. Here we employ an active learning approach that uses a data-driven model’s entropy as an uncertainty measure and explore relative to that entropy conditioned on the sensor state variables. Using a coverage-based ergodic controller, we train perceptual models in near-real time. We demonstrate our approach using a biomimetic sensor, exploring “tactile scenes” composed of shapes, textures, and objects. Each learned representation provides a perceptual sensor model for a particular tactile scene. Models trained on actively collected data outperform their randomly collected counterparts in real-time training tests. Additionally, we find that the resulting network entropy maps can be used to identify high salience portions of a tactile scene.

I. INTRODUCTION

Touch enables locating objects and navigating spaces without relying on vision. For instance, rummaging through a drawer or locating a light-switch in the dark are both parts of everyday life. The ability to operate in fully or partially vision denied environments, say inside a machine, is essential for a wide range of tasks in maintenance, inspection, and manufacturing. Over the last several decades, a number of sensors have been developed which provide human, or at least human-adjacent, levels of performance across a range of tactile sensing tasks. Despite significant work in the area, robust tactile sensing and perception is aspirational in most deployed robots today.

This gap between what we observe in tactile animal behavior and what we would like to observe in robotic systems is partially due to touch being an obligate active sense (in contrast to vision and aural senses that can be effective as passive sensors). The output of a sensor is dependent not only on the environment and object under test, but also on the contact conditions—e.g., relative velocity, acceleration, and pressure—of the sensor itself. Furthermore, biomimetic tactile sensors will, in almost all cases, be near-field and must explore an object extensively in order to collect information about its size and shape. Biomimetic tactile sensors are

¹Center for Robotics and Biosystems, Northwestern University, Evanston, IL, USA.

²Department of Mechanical Engineering and Materials Science, Yale University, New Haven, CT, USA.




Fig. 1. **Haptic exploration of a leaf on a test token:** The high entropy regions of the neural network indicate where the sensor should collect data in the scene—the image edge is the edge of the reachable scene—most relevant to predicting future measurement values. The black ellipse indicates the approximate size of the sensor footprint and the high information content areas include both when the sensor is in direct contact as well as when the sensor is adjacent to the leaf.

also frequently high dimensional, stochastic and non-linear, making manual processing of the data difficult. One solution to these challenges is to use data-driven (learned) perceptual models, that ingest raw data and return synthesized models of the tactile landscape.

We present a method for synthesizing exploratory behaviors to accelerate training of haptic perceptual models, using the SynTouch BioTac sensor as a model haptic sensor. Our primary contribution in this work is an active learning strategy for generating tactile exploratory motions by exploring relative to a generative model’s conditional uncertainty distribution. This enables the sensor to spend more time in high information areas of the space, as determined by the underlying machine learning model. We demonstrate that this method produces more accurate perceptual models during real-time training as compared to a random exploration baseline across six different tactile scenes with a variety of textures and geometries. We moreover demonstrate that the uncertainty maps from these models can be used to identify areas of high importance in a tactile scene, even when those areas are caused by organic objects like leaves which defy easy manual definition.

This paper is organized as follows: Section II discusses important background relating to the model architecture, sensor, and learning strategy employed in this paper. Section III

outlines the theoretical framework and software architecture used for our experiments. Section IV provides details about experimental hardware. Finally in section V we discuss the performance of our method.

II. RELATED WORK

This work relies on active learning and generative neural networks, in particular autoencoders. We also make use of the SynTouch BioTac to demonstrate our method. The following subsections provide additional background.

A. CVAEs

Autoencoders (AEs) are a neural network variant which find lower dimensional data representations in a manner analogous to non-linear principal component analysis. Autoencoders are typically composed of two elements: an encoder network and a decoder network. The encoder network takes full dimension input data and returns a lower dimensional latent representation. The decoder network takes an element of the latent space and returns a signal in the same space as the original input; training is based on requiring the decoder output to approximate the original input data. The encoder network can be used as a preprocessor for efficiently training other models to perform tasks such as classification or control.

To ensure better out-of-distribution performance, variational autoencoders (VAEs) modify the encoder to produce a multivariate latent distribution. This distribution is then sampled to provide a latent vector for the decoder network [1]. The latent distribution acts as a regularizer, and improves performance in regions of the latent space with sparse training data [2].

A final modification of the AE structure introduces a conditional vector associated with each data point. The conditional vector is provided as an input to both constituent networks as shown in figure 3. During inference the conditional input to the decoder can be varied to predict how sensor readings would change given different conditional parameters, so that the decoder can be used as a generative model. The resulting model is called a Conditional Variational Autoencoder (CVAE) [3].

Unsupervised training makes VAEs particularly promising for use in settings where novelty is expected (e.g., infrastructure monitoring [4]). VAEs and CVAEs have also been used for sensor modeling, including for compressing radar images and modeling soft sensors [5]–[7]. CVAEs are particularly promising for perceptual modeling, because the conditional vector provides a way to extract predicted sensor data across a variety of sensor states [8], [9], making them particularly appropriate for haptic perception.

B. Active Learning

Active learning (AL) is a machine learning strategy in which information about the model state or performance is used to guide data collection. This can result in lower losses or greater data efficiency relative to conventional training strategies. In deep learning AL can improve training times

by biasing the data set, or reduce data labeling costs by prioritizing certain samples for manual review [10]. In a robotics context AL can be used to guide the behavior of a system to maximize model learning. For example, an active learning approach increases the rate at which Koopman operators can be re-trained on simulated quadrotor dynamics in flight [11]. AL can additionally be powerful for training perceptual sensor models, since it enables a data collection system to focus on only the most important parts of a sensor target.

Training speed, data efficiency, and energy use are all key considerations when training learned sensor models. Active learning can be used to substantially improve training performance of CVAEs on camera-like sensors [12]. However, active learning has not been used for generating perceptual models of tactile sensors. Since tactile sensors tend to be delicate and slow to collect samples, they represent a particularly promising application area of AL, motivating the present work.

C. BioTac Sensor

We demonstrate our methods using a BioTac sensor from SynTouch. The BioTac is a biomimetic finger-tip sensor designed to have sensing capabilities similar to those of a human finger [13]. The BioTac has three primary sensing modalities: an ultrasonic pressure sensor which provides spectral data, an array of 19 electrodes which provide spatial pressure data, and a temperature sensor which provides a measurement of heat-flux [13]. The BioTac is composed of a solid inner core, containing the sensing electronics, a soft conductive skin, and an electrolytic fluid layer [13]. The outer skin contains “fingerprint” ridges which help the ultrasonic sensor resolve textures, and it is held in place by a rigid “fingernail” which also serves to seal the fluid chamber [13].

Early research using the BioTac showed near-human microvibration and impact detection performance using the ultrasonic sensor [14]. In [15] it was demonstrated that a set of 3 features generated from the ultrasonic data can be used to accurately identify various textures using Bayesian exploration. Recently, a spiking network was applied for rapid texture identification using the BioTac’s electrode array across 20 different materials [16]. The BioTac is extensively used in haptic robotic manipulation, including for haptic teleoperation and manipulation of diverse objects [17]–[20].

Although the BioTac sensor provides a rich and sensitive data stream, it experiences drift and is non-linear in force response beyond 2N [15]. Finite Element Analysis (FEA) simulations and data sets have improved modeling of the BioTac sensor [21]–[23]. However, while there are potentially adequate simulation capabilities for a BioTac in a manipulation context, there are no simulations suitable for predicting BioTac texture response.

III. METHOD

Our method efficiently trains tactile sensor models for processing tactile data to enable higher level robotic reasoning

(e.g., object identification, registration, and localization). We have selected a CVAE for the model architecture and an ergodic exploration for generating exploratory motions. We demonstrate this approach using a robotic gantry, a SynTouch Biotac sensor, and a number of 6"×6" tactile scenes.

A. Learning Architecture

Our learning process (Fig. 2) alternates between periods in which the model is trained using previously collected data, and periods in which the state of the model is used to guide the collection of new data.




Fig. 2. **Closed loop data collection for haptic learning:** After initially collecting data using a uniform distribution as the specification, the system collects data based on the state of the network while simultaneously training the network based on collected data.

B. Model Architecture




Fig. 3. **Network architecture:** The haptic measurement is conditioned on the sensor state while the decoder predicts the measurement at any sensor state.

Our work uses the CVAE architecture shown in Fig. 3. This model architecture was selected for its ability to efficiently compress structured data, and because it can be used to make predictions about sensor values for different system states. The model’s data input has 57 dimensions consisting of three concatenated 19 dimensional sensor readings, and its conditional input is the position of the sensor in Cartesian space. The model’s output distribution is a maximum likelihood Gaussian modeling the encoder input. The output distribution has a scalar diagonal variance matrix with a magnitude provided by the decoder. This variance output from the decoder provides a measurement of the network’s uncertainty for a particular pairing of a latent space element and conditional vector, which might reflect either underlying variance in the data or the training state of the network.

The encoder and decoder networks are composed of fully connected layers and taper symmetrically towards the center of the network. Each half of the network has an inter-layer ratio of 0.8, with an initial layer width of 300, and a depth of 4 layers. The encoder outputs means and variances for a 6 dimensional latent distribution with a diagonal variance matrix. During training a value is sampled from this distribution and passed to the decoder as the model’s latent vector. The model is trained with a batch size of 256 and a dropout value of $p = 0.2$. Leaky ReLU activation functions are used for all neurons except the output layer, which uses a sigmoid function to enforce consistent scaling. The network is initialized with Kaiming Uniform Initialization for $a = \text{sqrt}(5)$ [24].

C. Ergodic Metric

Definition 1 (Ergodic): An agent’s trajectory is ergodic with respect to some distribution, if the trajectory’s time-averaged statistics are equal to the distribution’s spatial statistics.

Generating exploratory motions for tactile sensing involves balancing two competing priorities: The sensor must explore high-salience parts of a scene under many contact conditions in order to build a repeatable model, but it must also continue to explore other parts of the environment and other contact conditions in case they contain high salience features. The coverage of a sensor trajectory relative to the environmental domain and contact conditions can be specified by minimizing the ergodic metric [25] relative to a continuous target distribution provided by a learning model.

In order to compare a trajectory (a set of states parameterized by one continuous variable time t) to a 2D (or higher dimensional) distribution, the trajectory is first decomposed into a series of delta functions, Eq. (1).

$$C(x) = \frac{1}{t} \int_0^t \delta(x - x(\tau)) d\tau \quad (1)$$

These delta functions can be represented using a Fourier decomposition of (1) using basis functions of the form (2) where n is the dimension of the state space, k is an index over the coefficients of the multidimensional Fourier transform, L_i is the length of the i^{th} dimension and h_k is a normalizing factor to ensure the basis remains orthogonal.

$$F_k(x(t)) = \frac{1}{h_k} \prod_{i=1}^n \cos\left(\frac{k_i \pi}{L_i} x_i(t)\right) \quad (2)$$

$$h_k = \left(\int_0^{L_1} \int_0^{L_2} \cos^2\left(\frac{k_1 \pi}{L_1} x_1\right) \cos^2\left(\frac{k_2 \pi}{L_2} x_2\right) dx_1 dx_2 \right)^{\frac{1}{2}} \quad (3)$$

The Fourier coefficients corresponding to the trajectory are given by (4) and the coefficients approximating the target distribution are given by (5).

$$c_k = \frac{1}{T} \int_0^T F_k(x(t)) dt \quad (4)$$




Fig. 4. **Sample scenes, exploration trajectory, network entropy:** For each data collection episode a sensor exploration trajectory is generated using the network entropy as a target distribution. This process ensures that the sensor spends time in informative areas of the scene.

$$\phi_k = \int_X \phi(x) F_k(x) dx \quad (5)$$

The ergodic metric can then be computed by taking the (Sobolev) distance between the time-averaged trajectory statistics and the target distribution (6).

$$\epsilon(t) = \sum_{k_1=0}^K \dots \sum_{k_n=0}^K (1 + \|k\|^2)^{-\frac{n+1}{2}} |c_k - \phi_k|^2 \quad (6)$$

For faster computation, we use a version of the sample-based ergodic controller described in [26]. This formulation approximates the ergodic measure based on Kullback-Leibler divergence and replaces the Fourier based ergodic metric above with (7). $P(s)$ is the target distribution, $q(s) = q(s|x(t))$, and s represents a sample.

$$D_{KL}(p||q) = -E_{p(s)}[\log(q(s))] \quad (7)$$

The expectation can then be replaced with a sample based approximation as shown in (8). In our case, the distribution being sampled is provided by the model under training. This allows the target distribution to evolve as the model learns and new areas of high tactile complexity are discovered.

$$D_{KL}(p||q) \approx - \sum_{i=1}^N P(s_i) \log(q(s_i)) \quad (8)$$

Haptic exploration benefits from this specification by always representing the entire domain and the likelihood that a sensor reading anywhere in the domain will improve the generative predictions for the haptic sensor. As a result, the exploration strategy can always take into account the entire domain *and* the state of the learning model over that whole domain, avoiding fixating on a small subset of the domain that happens to be feature rich.

D. Exploratory Motions

The premise of our approach to generating exploratory motions is that by spending more time in areas of the conditional space with high model entropy, the sensor will collect a higher quality data set. Since the model output is structured as a multivariate Gaussian, (9)—which describes the entropy of a multivariate Gaussian—can be used to calculate the model entropy for a given decoder input. For (9): $x \sim N_D(\mu, \Sigma)$ is the network output distribution, Y is the latent vector, Z is the conditional vector, and $D = 57$ is the dimension of a data point.

$$H(Y, Z) = \frac{D}{2}(1 + \log(2\pi)) + \frac{1}{2}\log(|\Sigma|) \quad (9)$$

Since for our model the output variance ($\Sigma = \sigma I_D$) is a scalar matrix, (9) simplifies to (10).

$$H(Y, Z) = \frac{D}{2}(1 + \log(2\pi)) + \frac{D}{2}\log(\sigma) \quad (10)$$

This value can then be sampled for different conditional vectors to provide the target distribution required for ergodic exploration. A fixed length trajectory is then optimized relative to this target distribution using the ergodic metric described by (8). As an example of this process, figure 4 shows two target distributions (Network Entropy) and the exploratory trajectories that result (Training Trajectory). In practice, each time the distribution is sampled we use several latent vectors drawn from the preceding round of data collection and then average the result. This ensures a consistent distribution during the trajectory generation process.

To encourage exploration we add a constant value to all sampled entropies, ensuring that at least 5% of the sensor's time is spent in low-entropy areas of the conditional space.

This helps prevent the neural network from fixating on high-entropy areas.

The initial data collection round—prior to any training—needs to be specified. To encourage exploration the initial target distribution is set to a uniform distribution across the conditional space. This ensures the sensor explores the entire scene before beginning the active learning process.

For comparison purposes, we use a passive sampling approach that does not use the state of the learning model, specifically relying on random walks, though raster scanning could also be used. Random walks can be vulnerable to over-exploring low information areas of the space, or to missing key areas—like a leaf, piece of leather, or another “object”—entirely. Rastering strategies ensure good coverage, but can take a long time to achieve good coverage and risk being dominated by irrelevant data. This is particularly an issue for tactile sensors, which must explore not only a range of positions/states, but also a range of contact parameters.

IV. HARDWARE

This work uses a SynTouch BioTac sensor mounted to a modified X-Carve gantry. Position tracking is provided by April Tags and a Logitech Brio. To ensure accuracy, the vision system and gantry are calibrated against each other at the start of each training run. The gantry speed is capped at $120 \frac{mm}{s}$ in order to minimize tracking error. All of the experiments in this paper were run using a desktop with a Ryzen 2970WX CPU and a GTX2080 GPU.

Six different tactile scenes were used in this work including various shapes cut from leather and acrylic, arrangements of painters and duct tape, and the leaves of a Zanzibar Gem plant. Each sample was arranged onto a 145mm x 145mm acrylic sample token as show in figure 5. These materials were chosen to represent a range of non-abrasive textures, and were arranged into various shapes to create “tactile scenes”. Where necessary samples were flattened and adhered to their acrylic carriers with double-stick tape, hot glue or CA glue.

The BioTac sensor is mounted on a pivoting compliant base with a preload of $1.25N$ and a spring constant of $0.39 \frac{N}{mm}$. Since the BioTac sensor is known to exhibit thermal drift, the contact force was re-calibrated prior to each data collection run. To further mitigate thermal drift the sensor was also allowed to run for at least an hour prior to use. However, we found that even with these measures the data varied substantially from day to day, and even between different collection runs on the same day. To protect the sensor’s skin from wear the pinkie finger of a small nitrile glove was used as an outer covering. This covering was replaced whenever any sign of wear occurred, similar to [15].

Data was collected from the BioTac at a rate of 100hz, and positions were recorded by the camera at a rate of 30hz. Only the BioTac’s 19 spatial electrodes were used—data from the heat flow and ultrasonic sensors were discarded. A single data point for training consisted of one position reading and the three preceding BioTac readings—resulting in 2 conditional dimensions and 57 electrode dimensions.

Prior to use each electrode value is normalized to a range of [0,1].




Fig. 5. **Experimental gantry, BioTac sensor and haptic tokens:** The interchangeable token each constitute a tactile scene and are held with pins for hot swapping. The sensor’s compliant mount can accommodate 6mm of height variation.

V. RESULTS

We tested our method using six different tactile scenes as described in section III. We selected these scenes to span a range of conditions including partial contact, multiple textures, complex shapes, and organic materials. For each scene we conducted two training runs using a random walk data collection strategy, and two runs using our active learning method. To provide a consistent comparison, we collected the same total path-length of data for all trials. As shown in figure 6 we found that the active learning method outperformed the baseline in every run on five of the six scenes. For the leather ‘N’, the random walk was expected to be comparable to the active learning approach since the high information areas cover almost the entire scene.




Fig. 6. **Average MSE losses for real-time training on 6 different tactile scenes:** Note that in all but one case the active ergodic exploration substantially outperformed the random walk. In the one case where the two are comparable, the Leather ‘N’, almost the entire scene has high entropy.

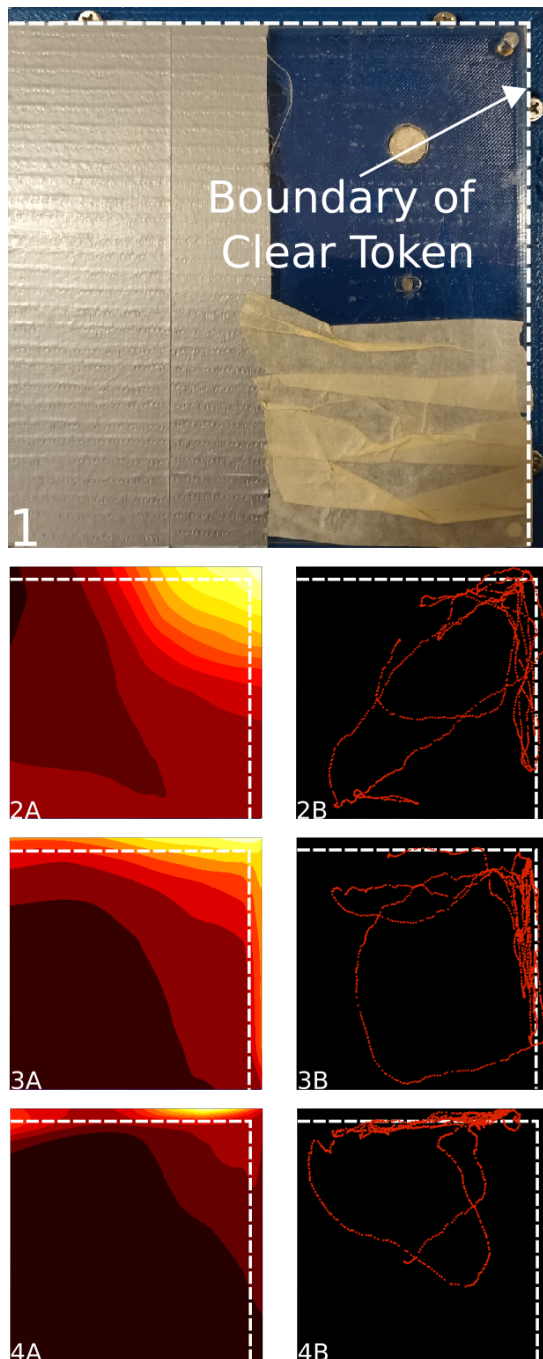


Fig. 7. **Haptic exploration naturally focuses on edges and corners:** When given access to the token edge (1), the model clearly identified the exposed edges and corner as high salience areas. The heatmaps (column A) show how the focus shifted along the boundary during different collection episodes. The trajectories in column B show the sensor movement for each episode.

Edges and corners are important haptic features, and during training the active exploration intermittently fixates on both edges and corners (Fig. 7). This behavior was most prevalent during a series of trials in which the system was able to access the upper edge of one of the tokens. These tests consistently produced entropy distributions similar to the one

in figure 7. That distribution clearly highlights the upper edge, right edge, and upper right corner of the tactile token. Moreover, the experimental response to detecting these edges and corners leads to the exploratory motion concentrating time on modeling them.

The natural emergence of this behavior is significant because edge detection is an important capability in tactile sensing. In human tactile sensing edges are salient features [27]. However, most current approaches for tactile edge finding either rely on bespoke detection conditions or on methods from computer vision like [28], [29]. These results indicate that our method is identifying some of the same high salience features as human touch without any explicit encoding of edges or corners in the exploration algorithm.

In addition to overall lower real time training losses, we also observed that the entropy heat maps generated by the active learning method were crisper and more refined than their passive learning counterparts. This is particularly notable, because we found that after network convergence, the entropy heat maps serve to highlight important “objects” in each tactile scene. This is particularly clear in Fig. 4 which shows entropy maps for both a leaf and the leather triangles. The placement and locations of high entropy areas are consistent between across training runs, suggesting that the entropy distribution itself can be used for object identification or localization.

VI. CONCLUSION

Tactile sensing requires contact and measurements will vary depending on how a sensor is articulated. Here we present a method for generating exploratory motions without relying on primitives, hand crafting, or specialized knowledge of the objects under test. Our method produces lower losses when training perceptual tactile sensor models than a random walk baseline. We use an active learning approach in which the sensor spends more time in areas of low neural network certainty and less time in areas of high network certainty, avoiding prescribing haptic perception in terms of pre-formed features. The network entropy distributions also serve as a salience map of the tactile environment—including organic, difficult-to-model scenes—potentially providing an opportunity for salience-based registration and localization.

We further show that without explicitly including edges and corners as haptic features, our method consistently finds edges and corners on simple scenes as an ‘emergent’ byproduct of connecting exploration strategy to the learning model. This type of edge-centric exploratory motion is important in human tactile sensing for object recognition and may enable the same tactile capabilities in robots. Future work includes using these haptic models for object identification, registration, and localization.

ACKNOWLEDGMENTS

J.K. and T.M. acknowledge support from the Army Research Office (ARO, Grant No. W911NF-22-1-0286).

REFERENCES

- [1] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," Dec. 2022.
- [2] A. Kumar and B. Poole, "On Implicit Regularization in VAEs," Dec. 2020.
- [3] K. Sohn, H. Lee, and X. Yan, "Learning Structured Output Representation using Deep Conditional Generative Models," in *Advances in Neural Information Processing Systems*, vol. 28, 2015.
- [4] Y. Zhang, X. Xie, H. Li, and B. Zhou, "An Unsupervised Tunnel Damage Identification Method Based on Convolutional Variational Auto-Encoder and Wavelet Packet Analysis," *Sensors*, vol. 22, no. 6, p. 2412, Jan. 2022, number: 6 Publisher: Multidisciplinary Digital Publishing Institute.
- [5] X. Zhu, S. K. Damarla, K. Hao, and B. Huang, "Parallel Interaction Spatiotemporal Constrained Variational Autoencoder for Soft Sensor Modeling," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 8, pp. 5190–5198, Aug. 2022.
- [6] J. Wang, S. Li, D. Cheng, L. Zhou, C. Chen, and W. Chen, "CVAE: An Efficient and Flexible Approach for Sparse Aperture ISAR Imaging," *IEEE Geoscience and Remote Sensing Letters*, vol. 20, pp. 1–5, 2023, conference Name: IEEE Geoscience and Remote Sensing Letters.
- [7] S. Dixit and N. K. Verma, "Intelligent Condition-Based Monitoring of Rotary Machines With Few Samples," *IEEE Sensors Journal*, vol. 20, no. 23, pp. 14 337–14 346, Dec. 2020.
- [8] X. Wang and H. Liu, "Data supplement for a soft sensor using a new generative model based on a variational autoencoder and Wasserstein GAN," *Journal of Process Control*, vol. 85, pp. 91–99, Jan. 2020.
- [9] M. Itkina, Y.-J. Mun, K. Driggs-Campbell, and M. J. Kochenderfer, "Multi-Agent Variational Occlusion Inference Using People as Sensors," in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 4585–4591.
- [10] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, "Active Learning with Statistical Models," Feb. 1996.
- [11] I. Abraham and T. D. Murphey, "Active Learning of Dynamics for Data-Driven Control Using Koopman Operators," Jun. 2019.
- [12] A. Prabhakar and T. Murphey, "Mechanical intelligence for learning embodied sensor-object relationships," *Nature Communications*, vol. 13, no. 1, p. 4108, Jul. 2022.
- [13] C. H. Lin, T. W. Erickson, J. A. Fishel, N. Wettels, and G. E. Loeb, "Signal processing and fabrication of a biomimetic tactile sensor array with thermal, force and microvibration modalities," in *2009 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, Dec. 2009, pp. 129–134.
- [14] J. A. Fishel and G. E. Loeb, "Sensing tactile microvibrations with the BioTac ; Comparison with human sensitivity," in *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. Rome, Italy: IEEE, Jun. 2012, pp. 1122–1127.
- [15] J. Fishel and G. Loeb, "Bayesian exploration for intelligent identification of textures," 2012.
- [16] T. Taunyazov, "Fast Texture Classification Using Tactile Neural Coding and Spiking Neural Network," *IEEE Xplore*, 2020.
- [17] J. Reinecke, A. Dietrich, F. Schmidt, and M. Chalon, "Experimental comparison of slip detection strategies by tactile sensing with the BioTac® on the DLR hand arm system," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 2742–2748.
- [18] J. Liang, A. Handa, K. V. Wyk, V. Makoviychuk, O. Kroemer, and D. Fox, "In-Hand Object Pose Tracking via Contact Feedback and GPU-Accelerated Robotic Simulation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 6203–6209, ISSN: 2577-087X.
- [19] J. A. Fishel, T. Oliver, M. Eichermueller, G. Barbieri, E. Fowler, T. Hartikainen, L. Moss, and R. Walker, "Tactile Telerobots for Dull, Dirty, Dangerous, and Inaccessible Tasks," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 11 305–11 310.
- [20] C. Pacchierotti, D. Prattichizzo, and K. J. Kuchenbecker, "Cutaneous Feedback of Fingertip Deformation and Vibration for Palpation in Robotic Surgery," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 2, pp. 278–287, Feb. 2016.
- [21] Y. S. Narang, B. Sundaralingam, K. Van Wyk, A. Mousavian, and D. Fox, "Interpreting and predicting tactile signals for the SynTouch BioTac," *The International Journal of Robotics Research*, vol. 40, no. 12-14, pp. 1467–1487, Dec. 2021.
- [22] P. Ruppel, Y. Jonetzko, M. Görner, N. Hendrich, and J. Zhang, in *Intelligent Autonomous Systems 15*, ser. Advances in Intelligent Systems and Computing, M. Strand, R. Dillmann, E. Menegatti, and S. Ghidoni, Eds., Cham, 2019, pp. 374–387.
- [23] Y. Chebotar, K. Hausman, Z. Su, A. Molchanov, O. Kroemer, G. Sukhatme, and S. Schaal, "BiGS: BioTac Grasp Stability Dataset," 2016.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, "Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification," Feb. 2015.
- [25] G. Mathew and I. Mezić, "Metrics for ergodicity and design of ergodic dynamics for multi-agent systems," *Physica D: Nonlinear Phenomena*, vol. 240, no. 4, pp. 432–442, Feb. 2011.
- [26] I. Abraham, A. Prabhakar, and T. D. Murphey, "An Ergodic Measure for Active Learning From Equilibrium," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 3, pp. 917–931, Jul. 2021.
- [27] M. A. Plaisier, W. M. Bergmann Tiest, and A. M. L. Kappers, "Salient features in 3-D haptic shape perception," *Attention, Perception, & Psychophysics*, vol. 71, no. 2, pp. 421–430, Feb. 2009.
- [28] J. Platkiewicz, H. Lipson, and V. Hayward, "Haptic Edge Detection Through Shear," *Scientific Reports*, vol. 6, no. 1, p. 23551, Mar. 2016.
- [29] N. F. Lepora, A. Church, C. de Kerckhove, R. Hadsell, and J. Lloyd, "From Pixels to Percepts: Highly Robust Edge Perception and Contour Following Using Deep Learning and an Optical Biomimetic Tactile Sensor," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 2101–2107, Apr. 2019.