Tactile Object Classification Using a Biomimetic Whisker Sensor Array and Spiking Neural Networks

Master thesis

T.J. van der Weij

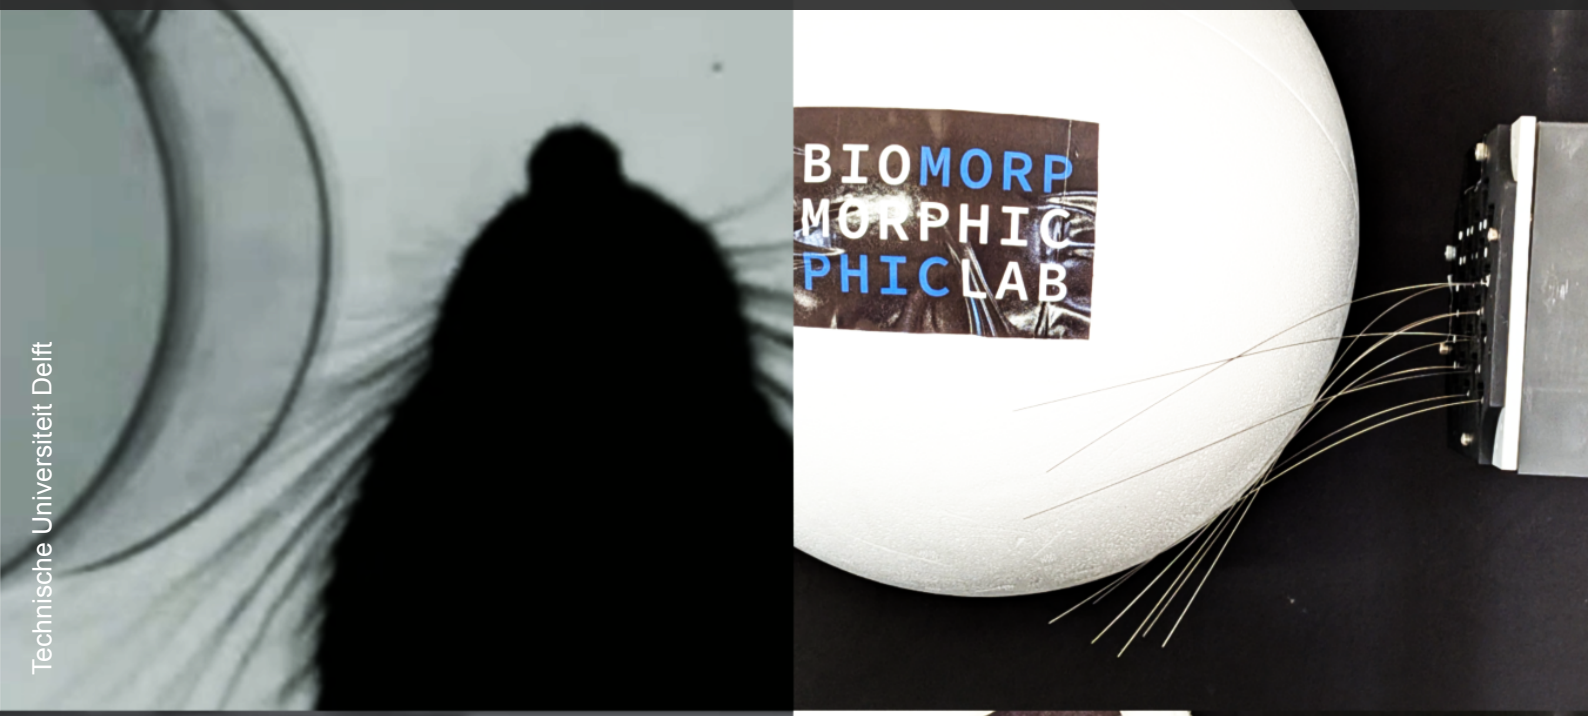

Technische Universiteit Delft

TU Delft

[This page is intentionally left blank]

Tactile Object Classification Using a Biomimetic Whisker Sensor Array and Spiking Neural Networks

by

T.J. van der Weij

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday April 22, 2026 at 12:30 PM.

Programme:	Robotics	
Department:	Cognitive Robotics	
Faculty:	Mechanical Engineering	
Student number:	4858999	
Project duration:	September 1, 2025 – April 22, 2026	
Supervisory Team:	Assoc. Prof. dr. ir. S. Hamaza,	TU Delft
	Asst. Prof. dr. N. Tömen,	TU Delft
	Assoc. Prof. dr. ir. L. Ferranti	TU Delft

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

[This page is intentionally left blank]

Preface

This master thesis has been conducted for the Master's program in Robotics, offered by the Department of Cognitive Robotics at TU Delft. It explores whisker-based tactile sensing and neuromorphic computing for robotics, drawing inspiration from nature.

This thesis project has been a long one: a year of diving into subjects spanning aerial tactile perception, whiskers (both biological and robotic), and spiking neural networks. It has been quite the ride, multidisciplinary at its core: from physical prototyping, 3D printing and metal working to software modelling, evaluation, and finally writing — a wide range of skills have been developed and improved along the way. I am proud of myself for staying true to my interests, remaining disciplined, and working passionately throughout the project.

Several researchers were involved in this project and made a great difference. I am grateful for the valuable support and feedback of my supervisors, Salua Hamaza and Nergis Tömen, both from TU Delft, throughout the thesis project. Their guidance has been greatly appreciated. I also wish to thank Laura Ferranti for her official supervision from the Department of Cognitive Robotics at TU Delft. I would also like to express my gratitude to the researchers who supported this work. Thanks to Teresa for providing the whisker sensor setup and for her continuous assistance with the experimental work and hardware. Thanks to Chaoxiang for providing additional context and insights into the whisker sensor and its functioning. Furthermore, Aurora's input was of great value, particularly regarding the theoretical background and practical implementation of neuromorphic computing. Finally, I would like to express my sincere gratitude to my family and friends, who have always supported me throughout my academic life.

T.J. van der Weij
Delft, April 2026

Contents

1

Introduction

Many animals navigate their surroundings efficiently by relying on touch rather than vision. Rats, harbour seals, and shrews, for example, sweep their whiskers across surfaces to detect obstacles, estimate object shape, and navigate in complete darkness. This tactile modality is fast, robust to visual occlusion, and provides direct feedback about contact geometry. Central to its efficiency is the brain's processing architecture: neural information is encoded as sparse, event-driven electrical impulses — spikes — through which complex, adaptive behaviour is achieved at remarkably low energy cost.

Robots are increasingly deployed in challenging, unstructured environments such as caves, tunnels, forests, and collapsed buildings. In such settings, conventional perception systems frequently fail: cameras require adequate lighting, LiDAR is degraded by dust and smoke, and GPS is unavailable underground. These limitations motivate alternative sensing modalities that remain functional during close-range physical interaction with the environment. Biomimetic whisker sensors address this need directly: passive and lightweight, they encode contact geometry through time-varying signals and operate independently of ambient lighting or atmospheric conditions.

Inspired by biological vibrissae, robotic whisker arrays have been developed for tasks including contour following, surface reconstruction, and object classification. Modern designs produce rich temporal signals during sweeping contact, moving well beyond binary collision detection. However, the processing pipelines used in these systems have typically relied on conventional neural networks, which can be computationally demanding on resource-constrained robotic platforms.

Spiking Neural Networks (SNNs) form a natural processing counterpart to whisker sensors: biologically motivated, they perform sparse, event-driven computation that mirrors neural processing and promises energy-efficient inference on neuromorphic hardware. Their temporal spike-based processing is inherently suited to time-varying sensor signals. Despite this complementary fit, no prior work has applied SNNs to the spatially distributed, time-varying pressure signals produced by a whisker array. This thesis addresses that gap by investigating whether a biomimetic whisker sensor array, combined with a fully spiking neural network, can classify objects from tactile time-series data — and under what conditions this approach succeeds.

The central question guiding this work is:

> *To what extent can a biomimetic whisker sensor array, combined with spiking neural networks, classify objects based on tactile time-series data?*

This is addressed along three themes — **classification performance**, **generalisation**, and **computational efficiency** — through the following sub-questions:

Classification performance

RQ-A Are raw pressure signals sufficient for classification, or do temporal derivative features improve performance?

RQ-D What is the minimum number of contact timesteps required for reliable object classification?

RQ-E How does the removal of individual whisker sensors affect classification performance?

Generalisation

RQ-B Does the trained SNN generalise to sweep speeds not encountered during training?

RQ-C Does the trained SNN generalise to the reversed sweep direction, and can joint training recover performance?

Computational efficiency

RQ-F How does the energy consumption of the SNN compare to a structurally equivalent non-spiking baseline?

The remainder of this report is organised as follows. **Chapter 2** presents the scientific paper, covering the full methodology, experimental setup, results, and conclusions. **Chapter 3** provides the literature study that contextualises and motivated this work. Supplementary material is included in the appendix.

2

Scientific Paper

Tactile Object Classification Using a Biomimetic Whisker Sensor Array and Spiking Neural Networks

T.J. van der Weij
Faculty of Mechanical Engineering, Delft University of Technology

***Abstract*—Robots operating in confined or visually degraded environments require sensing modalities that provide reliable close-range feedback. Biologically inspired whisker sensors address this need: passive and lightweight, they encode contact geometry through time-varying signals. Spiking neural networks (SNNs) form a natural processing counterpart — also biologically motivated, they perform sparse, event-driven computation that mirrors neural processing and promises energy-efficient inference on neuromorphic hardware. Despite this complementary fit, training data for neuromorphic tactile sensing remains scarce, and fully spiking architectures have not been evaluated on whisker-based classification, leaving their practical viability an open question.**

To address this gap, we collected two tactile datasets using a 10-whisker MEMS-based sensor array swept across three object geometries under varying contact conditions. A compact fully spiking network with 901 parameters achieved 100% test accuracy, outperforming larger spiking architectures; classification was reliable from as few as 20 timesteps of contact, and raw pressure signals alone were sufficient. Generalization experiments showed that sweep speed had minimal effect on performance, while indentation depth and sweep direction introduced larger domain shifts. Compared to a structurally equivalent non-spiking baseline, the spiking model matched classification accuracy with an estimated 9.5% reduction in overall inference energy under analog input and a potential 95% reduction under spike-encoded input.

These results demonstrate the feasibility of a fully bio-inspired pipeline — from biomimetic whisker sensing to spiking neural inference — for tactile object classification.

I. Introduction

Many animals navigate their surroundings efficiently by relying on touch rather than vision, using vibrissae (whiskers) to detect obstacles, estimate object shape, and navigate in complete darkness. Rats, harbor seals, and shrews, for example, achieve this by sweeping their whiskers back and forth across surfaces (Fig. 1, left) [1]–[3]. This tactile modality is fast, robust to visual occlusion, and provides direct feedback about contact geometry. Central to the efficiency of such biological systems is the brain's processing architecture: neural information is encoded as sparse, event-driven electrical impulses — spikes — the global currency through which complex, adaptive behavior is achieved at remarkably low energy cost. Imitating this natural intelligence in robotic systems is a central challenge of modern robotics.

Robots are increasingly deployed in complex, unstructured environments such as caves, tunnels, forests, and collapsed buildings. In such settings, conventional perception systems often fail: cameras require adequate lighting, LiDAR is degraded by dust and smoke, and satellite navigation signals

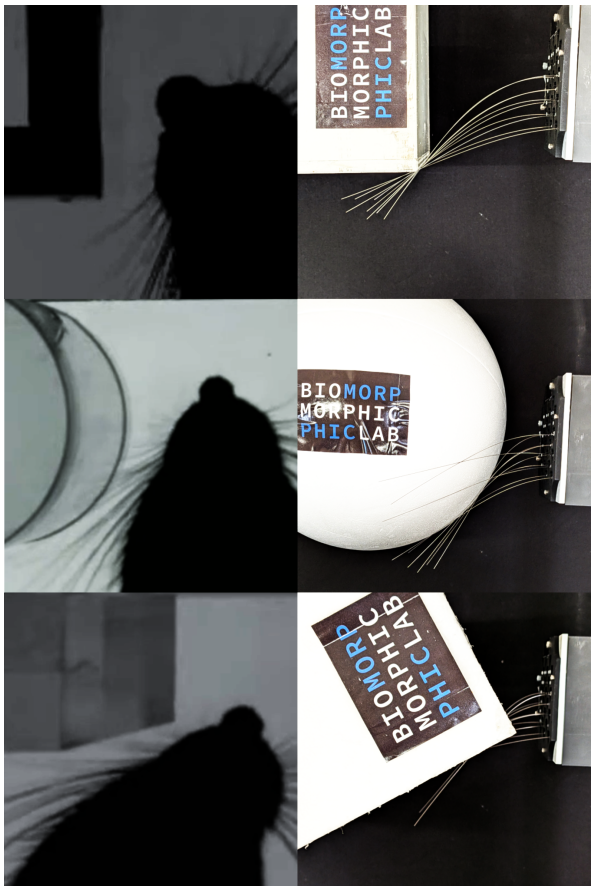

Fig. 1. Biological inspiration and robotic implementation. Left: a rat uses its vibrissae (whiskers) to explore objects through active whisking (image courtesy of the Active Touch Laboratory [4]). Right: the biomimetic whisker sensor array used in this study, sweeping across a test object in the laboratory.

are unavailable underground or indoors [5]–[7]. Even active acoustic sensing can be unreliable in cluttered or confined spaces. These limitations create a need for alternative sensing modalities that remain functional during close-range physical interaction with the environment. Relevant scenarios include autonomous navigation of caves and confined spaces [8], [9], search-and-rescue in collapsed buildings [10], covert operations benefiting from passive, low-power sensing [11], [12], and environmental monitoring in dense canopies or other environments inaccessible to camera-based systems [13].

Inspired by this principle, robotic whisker sensors have been developed to provide tactile feedback in environments

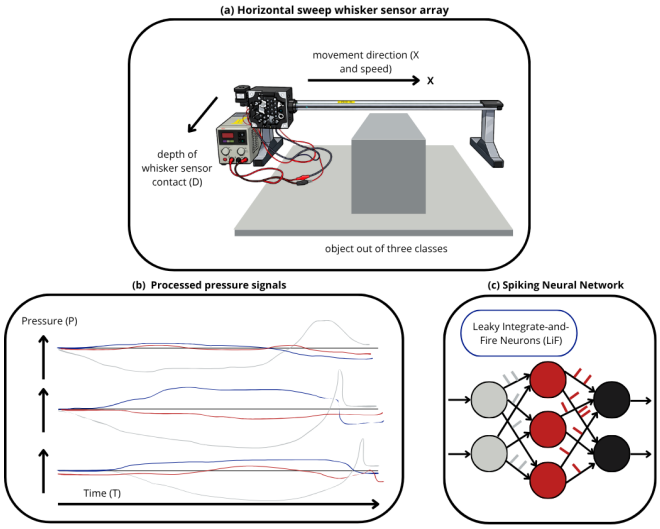

Fig. 2. Overview of the bio-inspired tactile classification system. (a) A 10-whisker MEMS-based sensor array is mounted on a linear stage and performs horizontal sweeps across three objects, recording three channel pressure signals for each whisker. (b) The resulting time-series data consist of 30 pressure channels capturing spatiotemporal tactile date as the array traverses the object surface. (c) The preprocessed pressure signals are fed as analog input to a spiking neural network, which integrates the temporal information through LIF neuron dynamics and classifies the trial into one of the three object classes via rate-coded output neurons.

where vision fails. Early designs used simple contact switches or bumpers to detect collisions and support wall-following behavior [3], [14]. While effective for basic obstacle avoidance, these systems provided only binary contact information. Modern whisker-inspired sensors, by contrast, produce time-varying signals during contact [15], encoding richer information about how and where an object interacts with the sensor (Fig. 1, right). This has enabled tasks such as contour following [16], local shape perception, object classification [17], surface reconstruction [18], and environmental mapping, including on aerial robotic platforms [5], [15], [19]. The key advance is that these sensors capture the full temporal dynamics of the whisker-object interaction, moving tactile sensing beyond collision detection toward precise perception of surfaces and objects. However, the processing pipelines used in these systems have typically relied on conventional neural networks, which can be computationally demanding for deployment on resource-constrained platforms.

At the same time, robots deployed in these environments face strict constraints on payload, power consumption, and onboard computation [20]. Tactile data must therefore be processed quickly and efficiently on embedded hardware. Conventional artificial neural networks (ANNs) can achieve high classification performance, but they are often computationally demanding for edge deployment [21]. Spiking neural networks (SNNs) offer a potential alternative: by processing information through sparse, event-driven activity rather than dense matrix operations, SNNs are widely considered a promising approach for low-latency and energy-efficient inference [22]. Moreover, the temporal spike-based processing in SNNs is naturally suited to time-varying sensor signals such as those produced by whisker arrays. This makes SNNs a relevant candidate for tactile perception on resource-constrained robots.

In parallel, neuromorphic tactile systems have demonstrated that SNNs can efficiently classify tactile data. For example, SNNs running on dedicated neuromorphic hardware have achieved high accuracy in texture recognition [23] and Braille reading [24], with substantial reductions in power consumption compared to conventional processing. These results suggest that SNNs are well suited to tactile time-series classification tasks. Yet, these studies have focused exclusively on fingertip-style sensors, and no prior work has applied SNNs to the spatially distributed, time-varying pressure signals produced by a whisker sensor array. Moreover, labeled datasets for neuromorphic tactile sensing remain scarce, further limiting progress in this area.

This gap is notable because whisker arrays and SNNs are a natural pairing: whisker sensors produce temporal pressure sequences during sweeping contact, and SNNs are inherently designed to process temporal spike-based information. Combining the two could yield a tactile perception system that is biologically inspired, fast and computationally efficient. This study addresses that gap by investigating whether a biomimetic whisker sensor array, combined with spiking neural networks, can classify objects based on tactile time-series data. Figure 2 shows the overall processing pipeline, from whisker-based pressure sensing to SNN-based classification.

This study investigates to what extent spiking neural networks can classify objects with different geometries using tactile pressure data from a biomimetic whisker sensor array. To address this, we pursue three objectives: **classification performance**, **generalization**, and **computational efficiency**.

To address the scarcity of training data for neuromorphic tactile sensing, we collected and contribute two labeled tactile datasets by sweeping a 10-whisker MEMS-based sensor array across three object geometries (cube, rectangular prism, and sphere) — chosen to represent typical shapes encountered in navigation and manipulation tasks — under varying indentation depths, sweep speeds, and motion directions. We then evaluated several fully spiking architectures — multilayer perceptrons, temporal convolutional networks, and recurrent SNNs — on this three-class classification task. Beyond overall accuracy, we investigated which input features, how many timesteps, and which whisker sensors were needed for reliable classification. We further tested whether the selected model generalized to unseen depths, speeds, and reversed sweep direction, and compared the energy consumption of the SNN with a structurally equivalent non-spiking baseline. All models were implemented using the snnTorch library [22].

The results showed that a compact three-layer spiking network with only 901 parameters achieved perfect classification on the held-out test set, outperforming more complex architectures. Generalization was robust to changes in sweep speed but degraded under larger domain shifts such as reversed sweep direction; the SNN matched the non-spiking baseline in accuracy while offering estimated energy savings of up to 95% under spike-encoded input.

These findings demonstrate that whisker-based tactile object classification using SNNs is feasible, and that spiking models

can match conventional network performance while offering a path toward substantially more efficient inference — making them a promising candidate for tactile perception on resource-constrained robots operating in environments where conventional sensing fails. The main contributions of this work are as follows:

- We contribute two labeled tactile datasets for whisker-based tactile object classification, collected under controlled conditions with a 10-whisker MEMS-based sensor array across varying depths, speeds, and sweep directions.
- We present a proof-of-concept fully bio-inspired pipeline combining biomimetic whisker sensing with spiking neural inference, and evaluate multiple fully spiking architectures on this task, analyzing robustness under varying contact conditions including sweep speed, indentation depth, and motion direction.
- We compare SNN models with non-spiking ANN counterparts and analyze their relative computational efficiency, highlighting the potential suitability of SNNs for resource-constrained robotic platforms.

II. Related Work

A. Robotic Whisker Sensing

Early robotic whisker systems used contact switches or geometric algorithms for collision detection and surface profiling [3], [25], but provided only limited temporal resolution.

Modern whisker sensors produce richer, time-varying signals during contact. MEMS-based designs have enabled lightweight, high-sensitivity whiskers suitable for aerial platforms [15], [26], while vision-based approaches use cameras to track whisker deflections optically [27], [28]. Multi-element arrays have also been developed for simultaneous identification and localization of objects [29], [30]. These advances have moved whisker sensing beyond binary contact detection toward continuous tactile perception, where time-series contact signals can encode object shape, texture, and contact geometry.

B. Whisker-Based Classification and Mapping

Whisker sensors have been applied to a range of tactile tasks. Kim et al. [31] demonstrated that biomimetic whiskers could discriminate object shapes based on deflection signals during active whisking. More recent work has combined whisker arrays with deep neural networks for object classification from point clouds [17] and with Bayesian filtering for 3D surface reconstruction [18]. Tactile navigation has also been explored, from simple surface identification on mobile robots [32] to aerial contour following [16] and quadrotor navigation using a compliant tactile finger [5].

A common theme across this work is that richer sensor signals — capturing the full temporal dynamics of whisker-object interaction rather than only instantaneous contact — enable more informative perception. However, the processing of these time-varying signals has relied on conventional neural networks or geometric methods. No prior work has applied spiking neural networks to whisker sensor data.

C. Spiking Neural Networks for Tactile and Time-Series Data

Unlike conventional ANNs, which propagate continuous-valued activations through dense matrix operations, SNNs communicate through sparse, binary spike events, making them naturally suited to neuromorphic hardware. SNNs process information through these discrete spikes over time, offering potential advantages in energy efficiency and temporal data processing [22], [33]. For time-series tasks, Lv et al. [34] demonstrated that SNN counterparts of TCNs, RNNs, and Transformers achieved comparable forecasting performance to conventional ANNs with substantially lower estimated energy consumption.

In tactile sensing specifically, neuromorphic systems have combined event-based sensors with SNNs. Brayshaw et al. [23] achieved 94% texture classification accuracy using a neuromorphic tactile sensor and SNN on Loihi2 hardware, with 96% power reduction compared to GPU execution. Xu et al. [24] applied spiking convolutional networks to Braille reading with a neuromorphic tactile sensor, demonstrating generalization across tapping speed and sensor position. Ward-Cherrier et al. [35] introduced the NeuroTac sensor, which transduces contact information into spike trains using an event-based camera.

These studies demonstrate the viability of SNNs for tactile classification tasks. However, they have focused on fingertip-style sensors rather than whisker arrays, and the sensing modalities differ fundamentally: fingertip sensors measure local deformation at the point of contact, while whisker arrays capture spatially distributed, time-varying contact signals.

III. Experimental Setup

A. Whisker Sensor

Each individual whisker sensor unit consists of three components [15]: (1) a nitinol wire serving as the whisker shaft, (2) a follicle structure containing a plastic tube, UV resin, and three MEMS barometers, and (3) an integrated microcontroller PCB for signal acquisition and communication. The whisker shaft is a 200 mm nitinol wire with a 0.4 mm diameter, which provides broad sensing coverage and high sensitivity while keeping stiffness low.

Upon contact with an object, forces and moments are transmitted through the whisker shaft to the follicle, where the resulting deformation exerts pressure on the embedded barometers. The whisker base experiences three forces and three moments, effectively encompassing all mechanical information entering the follicle [15]. The three BMP390 barometers (Bosch) at the whisker base measure the resulting pressure changes, providing a three-dimensional pressure response. Bending the whisker toward a barometer compresses the surrounding resin, producing a negative pressure reading, while bending away produces a positive reading. This directional sensitivity enables the sensor to capture the full spatial structure of the contact.

Signal acquisition is managed by an STM32F070F6 microcontroller, which reads the three barometers sequentially via an SPI bus. Each hexagonal PCB features communication ports on all sides, enabling flexible daisy-chain expansion

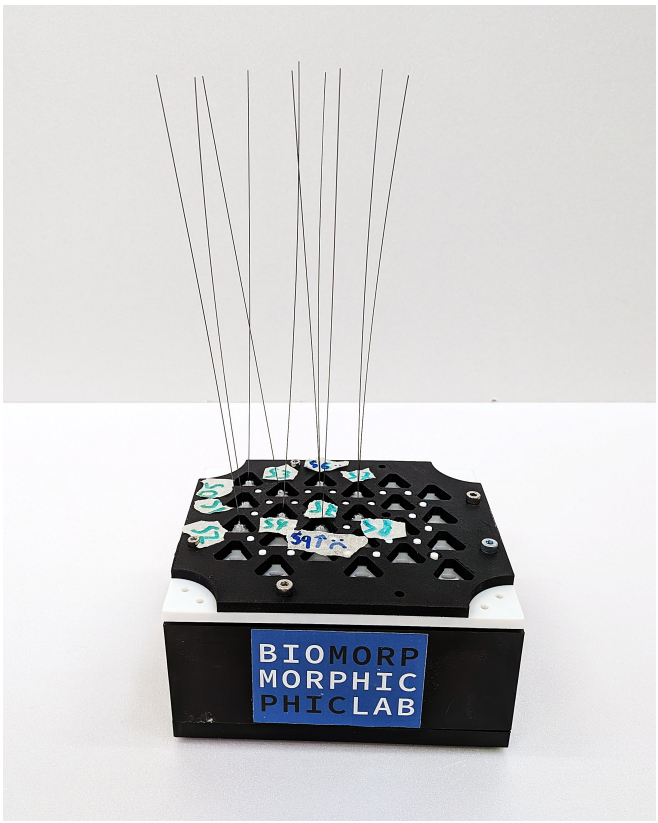

Fig. 3. Whisker sensor array with ten whisker units arranged in four columns (top to bottom, left to right: 0 1 2; 3 4; 6 5 9; 7 8). During a forward sweep the top row makes contact first; during a backward sweep the bottom row makes contact first. Sensors 6 and 9 (marked) were excluded due to non-functional or low-sensitivity readings, leaving eight active whiskers (24 pressure channels).

into larger arrays. However, the use of soft material-based MEMS whiskers inherently introduces limitations such as non-linearity and hysteresis in the sensor response [15].

B. Whisker Sensor Array

Ten whisker units of the design described above are arranged in a fixed spatial layout and connected in a daisy-chain configuration through the hexagonal PCB communication ports. The resulting array records 30 pressure channels in total (10 whiskers $\times$ 3 barometers each). An analog multiplexer (Analog MUX Click, MIKROE-4111) is used to route the sensor signals, and a Seeeduino XIAO nRF52 microcontroller coordinates readout across all units.

The physical layout of the array is shown in Fig. 3. During initial characterization, two whisker units (sensors 6 and 9) were found to be non-functional or insufficiently sensitive. Their corresponding six pressure channels were excluded during preprocessing, leaving 24 input channels (8 whiskers $\times$ 3 channels) for all modeling experiments.

C. Experimental Platform and Test Objects

The experimental platform consists of the whisker sensor array mounted on a motorized linear stage that performs controlled sweeping motions across test objects.

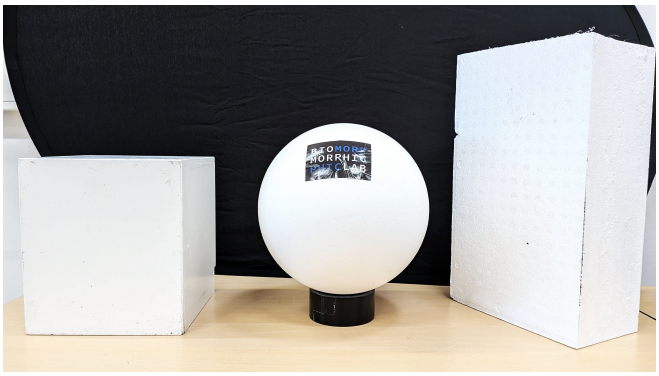

Fig. 4. The three test objects used for classification: sphere, cube, and rectangular prism, oriented at an angle to present a corner-like contact profile.

The array is attached to a belt-driven linear stage (855 mm travel) actuated by a NEMA17 stepper motor. A Seeeduino XIAO nRF52 microcontroller simultaneously controlled the stepper motor and read all whisker sensors via the analog multiplexer, logging synchronized pressure and motor-position data to the host computer.

Three rigid objects with distinct contact geometries are used for evaluation: a sphere (diameter 300 mm), a cube (250 mm side length), and a rectangular prism ($280 \times 160 \times 470$ mm), oriented at an angle to present a corner-like profile. During each trial, the whisker array sweeps toward the object at a constant velocity. The motion consists of four phases: approach, contact, dwell, and retraction, covering a total stroke of 700 mm. Recording begins from a fixed zero position located 250 mm from the object surface.

Motor steps are converted to millimeter displacement using the known belt pitch, providing a position estimate at each time step. Contact depth is defined as the perpendicular distance between the whisker tips and the object surface at the moment of first contact. Pressure signals are sampled at approximately 5 Hz, producing a frame-based dataset. To reduce experimental bias, the order of object type, sweep speed, and contact depth is randomized across trials.

D. Tactile Dataset Acquisition

Two datasets were collected using the experimental setup described above. Representative pressure signals and preprocessing visualizations for both datasets are presented in Section VI-A.

1) Forward-sweep dataset: The first dataset consists of forward sweeps (defined as motion toward the object) against the three object classes. In each trial, the whisker array approaches the object, makes contact, and continues the sweep along the slide axis while pressure signals are recorded.

The forward-sweep dataset varies along three experimental factors:

1) Object type: cube, rectangular prism, sphere
2) Indentation depth: 55 mm, 80 mm, 100 mm
3) Lateral sweep speed: 19 mm/s, 26 mm/s, 38 mm/s

Five trials were recorded for each combination of conditions, resulting in a total of $3 \times 3 \times 3 \times 5 = 135$ trials.

During each trial, synchronized measurements include the 24 active pressure channels, the motor position along the slide axis, and the corresponding labels (object type, sweep speed, and indentation depth). Because sweep speed varies across conditions, trial duration and sequence length also vary.

2) Backward-sweep dataset: A second dataset was collected to evaluate generalization across motion direction. In this dataset, the whisker array performs a backward sweep (defined as motion away from the object) after initial contact.

To isolate the effect of sweep direction, the motion parameters were fixed to a single configuration:

1) Indentation depth: D_2 (80 mm)
2) Lateral speed: S_3 (38 mm/s)

Thirty backward-sweep trials were collected in total, with ten trials per object class. Data were recorded using the same sensing and acquisition setup as for the forward-sweep dataset.

IV. Data Processing

A. Preprocessing

Raw trial recordings were transformed into model-ready time-series tensors using a fixed preprocessing pipeline consisting of contact-phase cropping, outlier handling, baseline correction, filtering, normalization, and sequence collation.

The signals were first cropped to retain only the contact-relevant portion of each trial. Timesteps corresponding to the initial approach phase, during which the sensor moved through air without object contact, were removed. For forward sweeps, this non-contact region corresponded to whisker-array positions from 0 to -250 mm. For backward sweeps, the non-contact range from 0 to $+250$ mm was removed. This cropping step ensured that the retained sequences focused on the physically meaningful contact phase.

To improve signal quality, extreme values exceeding 200 kPa were replaced with NaN, followed by IQR-based clipping with factor 3.0.

After outlier handling, baseline correction was applied independently to each channel by subtracting the first valid value. The resulting signals were then filtered using a first-order Butterworth band-pass filter (0.03–2.0 Hz at $\approx$5.16 Hz sampling rate), removing slow drift and high-frequency noise. Finally, the processed signals were standardized using global z-score normalization computed on the training split only. Because trial duration varied across sweep conditions, sequences were padded to a common length with a binary mask to exclude padded timesteps from loss and prediction.

B. SNN Model Architectures

Three fully spiking model families were evaluated: multilayer perceptrons (MLPs), temporal convolutional networks (TCNs), and recurrent neural networks (RNNs). Each architecture follows a standard ANN design, with activation functions replaced by leaky integrate-and-fire (LIF) neurons to form a fully spiking equivalent. This approach allows direct structural comparison between spiking and non-spiking variants while leveraging architectures with well-understood inductive biases for time-series data. These families represent the most widely used architectures for deep learning-based time-series classification [36]: MLPs provide a simple, parameter-efficient baseline with no structural assumptions about the input; TCNs exploit translation invariance along the time axis, enabling detection of contact patterns regardless of their temporal offset within the sweep; and RNNs incorporate explicit recurrent memory, suited to signals where the history of contact shapes the current response. Because the dataset is relatively small, all models were intentionally kept low-dimensional to reduce the risk of overfitting.

1) Fully Connected Models: Two MLP variants were evaluated: MLP3 and MLP4. Both consisted of fully connected layers interleaved with leaky integrate-and-fire (LIF) neurons. The input layer received the 24-dimensional tactile feature vector at each timestep, and the final layer produced three output neurons corresponding to the object classes. MLP3 contained two hidden layers, while MLP4 contained three. The hidden layer size was varied between 32 and 64 neurons during the hyperparameter sweep.

2) Temporal Convolutional Models: Temporal convolutional networks (TCNs) were used to capture local temporal patterns in the tactile signals. These models applied one-dimensional convolutions over the time axis while treating the sensor channels as input channels. The architecture consisted of two Conv1d layers, each followed by LIF neurons, with optional temporal max pooling ($k = 1$ or $k = 2$) along the time axis before a fully connected spiking classifier.

Two variants were considered. The standard TCN mixed all sensor channels in the first convolutional layer. The TCN_Whisker variant used grouped convolution in the first layer, where the three channels belonging to each whisker were processed independently before cross-whisker integration occurred in the second convolutional layer. This design reflected the physical grouping of the tactile sensors.

3) Recurrent Models: A recurrent SNN architecture was evaluated using recurrent leaky integrate-and-fire neurons (RLeaky) from snnTorch. In this neuron model, each neuron's previous spike output was fed back as additional input current, giving it an explicit memory of its own recent firing history beyond the standard membrane decay.

Two configurations were considered: a single recurrent hidden layer and a two-layer recurrent architecture. The output layer could optionally use recurrent connections, allowing evaluation of whether recurrence in the classifier stage improved performance.

C. ANN Baseline Model

To isolate the effect of spiking dynamics on classification performance and energy consumption, a non-spiking ANN counterpart was constructed for the best-performing SNN architecture identified during the hyperparameter sweep (Section VI-C). This ANN model, referred to as MLP3_NonSpike, replicated the connectivity of the spiking MLP3 two fully connected layers with identical layer sizes (24→32→3) and no dropout but replaced the LIF neurons with ReLU activations and removed all membrane dynamics. Instead of spike-count decoding, the ANN produced class predictions

by mean-pooling the per-timestep output logits over the full sequence length. This ensured that differences in classification performance and computational cost could be attributed to the spiking mechanism rather than to architectural differences.

D. Training Procedure

Fully spiking neural networks were trained to perform three-class object classification from complete whisker-pressure sequences. Each sample corresponded to one full trial and was assigned a single object label. Unless stated otherwise, training on the forward-sweep dataset pooled data across all sweep speeds and indentation depths.

All models used the standard leaky integrate-and-fire (LIF) neuron update rule with a fixed threshold of 1.0, trained via backpropagation through time using surrogate gradients (arctangent) [22]. Normalized raw sensor values were provided directly as analog input currents to the first spiking layer, avoiding the information bottleneck of external rate encoding. The optimizer was selected as part of the hyperparameter sweep (Section V-C), with Adam and AdamW producing the most reliable convergence. Training ran for a maximum of 200 epochs with early stopping based on validation performance (patience of 50 epochs).

Class prediction was based on spike-rate decoding. Output spikes were accumulated across time, and the predicted class was taken as the output neuron with the highest total spike count over the sequence. Training used a cross-entropy spike-rate loss computed from the output spike activity over time. Variable-length sequences were padded and masked as described in Section IV-A.

V. Experimental Protocol

A. Data Splitting and Evaluation

All experiments were conducted at the trial level, where each recorded trial was treated as one independent sample. To prevent leakage between temporally related measurements from the same recording, splitting was performed over complete trials rather than over individual time steps or subsequences.

For the main i.i.d. classification experiments on the forward-sweep dataset, the data were divided into training, validation, and test sets using a stratified split. Stratification was performed over object class to preserve class balance across the three subsets. The split was generated once using a fixed random seed of 42 and stored as a JSON file so that all models were evaluated on exactly the same partition. For the 135-trial forward-sweep dataset, this resulted in 93 training trials, 18 validation trials, and 24 test trials.

The validation set was used for all model-development decisions; the test set was held out until after all design choices were finalised.

Model performance was assessed primarily using accuracy and macro-F1. For the domain generalization experiments, results were additionally reported per held-out domain, and confusion matrices were used to analyze class-specific error patterns. Unless stated otherwise, each reported condition was evaluated over three random seeds, and results were reported as mean $\pm$ standard deviation.

TABLE I
Preprocessing ablation conditions.

Condition	Description	Purpose
C1_N0	Cropped + min–max	Baseline
C0_N0	Uncropped + min–max	Effect of cropping
C1_N1	Cropped + z-score	Effect of normalization

C0/C1: uncropped/cropped to the contact phase (position ≤ -250 mm); N0: global min–max scaling to $[0, 1]$; N1: global z-score ($\mu = 0$, $\sigma = 1$). Each row varies one factor relative to the baseline (top row). All other settings held fixed: MLP3, raw analog input, Adam, batch size 16, 200 epochs with early stopping.

B. Fixed Pipeline

Using the preprocessing operations defined in Section IV-A, a controlled ablation was conducted to select the default cropping and normalization strategy. The purpose of this phase was to determine one fixed preprocessing pipeline for all subsequent experiments.

All preprocessing ablations were conducted using a fixed baseline configuration: three-class object classification with MLP3, raw analog input, spike-rate decoding, cross-entropy spike-rate loss, Adam optimization, batch size 16, and a maximum of 200 epochs with early stopping (patience = 50). Each condition was evaluated over three random seeds on the validation set using accuracy and macro-F1.

The ablation focused on two factors: cropping and normalization. Cropping determined whether the sequence was restricted to the contact-relevant portion of the sweep, while normalization defined how pressure values were scaled before training. The evaluated conditions are listed in Table I. The filtering configuration was kept fixed, as the selected band-pass range of 0.03–2.0 Hz was treated as a measurement-driven design choice rather than an experimental variable.

C. Hyperparameter Sweep

After establishing the fixed preprocessing pipeline, a constrained hyperparameter sweep was conducted to identify the best-performing fully spiking architecture for the object-classification task. The search focused on architectural and training hyperparameters while keeping preprocessing fixed. The architectural search space is summarized in Table II, and the neuron and optimizer parameters in Table III.

The sweep was performed using random search with 150 trials. Random search was chosen because the search space contained many categorical and discrete parameters across different model families, for which it provided broad and uniform coverage.

The candidate model families were MLP3, MLP4, TCN, TCN_Whisker, and RNN. All models used the fixed analog-input pipeline with global z-score normalization. Performance was evaluated on the validation split using macro-F1, and the best configuration was selected based on validation performance. To reduce variance due to stochastic optimization, each sampled configuration was evaluated over three random seeds and the mean validation score was used for selection.

Because the validation set contained only 18 trials, the sweep was intentionally kept constrained to reduce the risk

TABLE II
ARCHITECTURE HYPERPARAMETER SEARCH SPACE.

Parameter	Values	Applies to
num_hidden	32, 64, 128	All models
ch1, ch2	16/32, 32/64	TCN, TCN_Whisker
kernel_size	3, 5	TCN, TCN_Whisker
pool_k	1, 2	TCN, TCN_Whisker
num_layers	1, 2	RNN
output_recurrent	true, false	RNN

num_hidden: hidden layer width; ch1/ch2: first/second Conv1d output channels; pool_k: temporal max-pooling kernel ($k=1$: no pooling, $k=2$: halves temporal dimension); output_recurrent: whether the RNN output layer uses recurrent (RLeaky) or non-recurrent (Leaky) neurons.

TABLE III
NEURON AND OPTIMIZER HYPERPARAMETER SEARCH SPACE.

Parameter	Search space
β	uniform $[0.5, 0.99]$
learn_beta	true, false
optimizer	Adam, AdamW, SGD
learning rate	log-uniform $[10^{-4}, 10^{-2}]$
weight decay	log-uniform $[10^{-5}, 10^{-2}]$
batch size	8, 16, 32

β: LIF membrane decay factor (higher values retain more membrane potential across timesteps); learn_beta: whether β is optimized during training or kept fixed.

of overfitting to the validation split. In addition to validation performance, the number of model parameters was considered as a selection criterion, with preference given to smaller models to reduce overfitting risk on the small dataset. After the sweep, the selected configuration was locked and evaluated once on the held-out test set.

D. Research Question Experiments

After the preprocessing pipeline and baseline model configuration were fixed through the preprocessing ablation and constrained hyperparameter sweep, the research questions were evaluated through controlled ablation experiments. These experiments address two of the objectives introduced in Section I: *classification performance* (effect of finite difference features, temporal window, whisker ablation) and *generalization* (depth and speed, direction). In each experiment, the preprocessing pipeline, selected architecture, and tuned hyperparameters were kept fixed, and only the factor under investigation was varied. All conditions were trained from scratch, and the held-out test set was evaluated once per condition per seed.

1) Effect of finite difference features: Sudden changes in pressure — such as those occurring at object edges or during whisker release — may carry discriminative information that is not equally apparent in the raw pressure magnitude. We therefore hypothesized that including finite difference features alongside raw pressure would improve classification. To evaluate this, only the input representation was varied. Three conditions were considered: raw pressure only (D0), finite difference only (D1), and raw and finite difference features concatenated (D2). The first-order finite difference was computed at runtime as

$$\Delta X[t] = X[t] - X[t-1],$$

where $X[t]$ denotes the z-score normalized pressure signal at time step t, and $\Delta X[0] = 0$. Finite difference features were normalized using statistics computed on the training split only. Under this setup, D0 and D1 used 24 input channels, whereas D2 used 48 input channels. All other model and training settings remained unchanged.

2) Temporal window: This experiment investigates how much temporal information is required for reliable object classification. Specifically, we evaluated at which point during the whisker sweep the SNN could make an accurate prediction.

To ensure consistent temporal alignment across samples, only trials at a fixed sweep speed (S1) are used. This subset contains 45 trials (15 per class), with nearly uniform sequence lengths. As a result, timestep index N corresponds to the same absolute time window across all trials.

Two types of truncated input sequences were evaluated:

- **Head (H):** the first N timesteps, capturing contact onset dynamics
- **Tail (T):** the last N timesteps, capturing late-contact or full-engagement dynamics

Comparing head and tail conditions at the same N allows us to determine whether early or late parts of the signal contain the most discriminative information.

The dataset is split at the trial level into 30 training, 6 validation, and 9 test samples.

A range of temporal windows was evaluated by varying $N \in \{10, 20, 30, 40, 50, 60, 80\}$ timesteps, corresponding to approximately 1.9 to 15.5 seconds of data. In addition, a full-length condition ($N = 116$) is included as a reference. Both head and tail truncations are tested for each value of N, resulting in a total of 15 experimental conditions.

Each condition was evaluated using three random seeds, yielding 45 runs in total.

3) Whisker ablation: This experiment evaluates the contribution of individual whiskers to the classification task. The baseline MLP3 model uses all 24 input channels (8 whiskers with 3 pressure channels each). Here, we assess whether reliable classification can be achieved using only a single whisker and identify which whiskers carry the most discriminative information.

Two types of input configurations were considered. In the full-input condition (S_all), all whiskers are used. In the single-whisker conditions (S0 to S7), only one whisker is provided as input, corresponding to its three pressure channels. This results in nine experimental conditions in total.

For the single-whisker settings, the MLP3 architecture is adapted by reducing the input dimensionality to three channels. Apart from this change in the first layer, all model architecture and training parameters remain identical to the baseline configuration. This ensures that performance differences can be attributed solely to the available sensor information.

Each condition was evaluated using three random seeds, resulting in 27 runs in total.

4) Depth and speed generalization: To evaluate robustness to unseen contact conditions, domain generalization experiments were conducted across indentation depth and sweep speed. In contrast to the random train/validation/test split used in the hyperparameter sweep, these experiments used domain-based splits over the full forward-sweep dataset of 135 trials. Although these splits overlap with the trials used in earlier phases, this did not introduce leakage because all models were trained from scratch within each domain generalization fold.

In the **leave-one-depth-out** (LODO) setting, models were trained on trials from two indentation depths and evaluated on the held-out depth. This yields three folds:

$$d_1 + d_2 \to d_3, \quad d_1 + d_3 \to d_2, \quad d_2 + d_3 \to d_1.$$

In the **leave-one-speed-out** (LOSO) setting, the same procedure is applied across sweep speed, again resulting in three folds:

$$s_1 + s_2 \to s_3, \quad s_1 + s_3 \to s_2, \quad s_2 + s_3 \to s_1.$$

Within each fold, two training conditions were compared: no augmentation (A0) and additive Gaussian noise jitter [37] during training (A1). This noise models realistic session-to-session sensor variability: barometric sensors can respond slightly differently across sessions due to temperature drift, mechanical settling, and contact repeatability. The noise was injected into the raw baseline signal before band-pass filtering, so that any artificially introduced high-frequency components were attenuated by the same filter applied to the original data. Appendix A, Fig. 15, illustrate the magnitude and qualitative effect of the selected augmentation level. Augmented samples were generated only for training trials; validation and test trials remained unaugmented. Under the balanced design, each held-out depth or speed domain contained approximately 45 trials, corresponding to 15 trials per object class. Performance was reported as accuracy on each held-out domain, together with confusion matrices.

5) Direction generalization: In real deployment a robot may approach objects from any direction. This experiment evaluates whether a model trained on forward sweeps generalizes to backward sweeps. The forward dataset was restricted to depth D_2 and pooled across all three speeds, resulting in 45 trials. The backward dataset consisted of 30 trials collected at the same depth D_2 and fixed speed S_3.

Three conditions were evaluated. In C1, only backward data were used, providing a reference for whether the backward dataset is intrinsically classifiable. In C2, the model was trained on forward data only and tested on backward data, forming the zero-shot transfer setting. In C3, the model was trained jointly on forward and backward data to test whether exposure to both directions reduces the transfer gap.

For all direction-generalization conditions, z-score normalization statistics were computed on the corresponding training split only. In the zero-shot setting, normalization parameters were fitted on the forward training set and applied unchanged to the backward test set. Each condition was evaluated over three random seeds using validation and test accuracy and macro-F1.

TABLE IV
PER-LAYER OPERATION COUNTS FOR THE ANN AND SNN MLP3 MODELS.

Layer	ANN	SNN (analog)
fc1 (24→32)	$T \times 768$ MACs	$T \times 768$ MACs
LIF updates	—	$T \times 35$ ACs
fc2 (32→3)	$T \times 96$ MACs	$r_{\text{LIF}} \cdot T \times 96$ ACs

(T = number of timesteps, r_{LIF} = hidden firing rate) MAC: multiply-accumulate (dense, used by ANN); AC: accumulate-only (sparse, triggered by spikes in SNN). LIF updates account for one AC per neuron per timestep for membrane potential decay. fc1 costs are identical because both models receive analog input.

E. Comparison with ANN Models

This section addresses the third objective introduced in Section I: *computational efficiency*. To isolate the effect of spiking dynamics, the non-spiking ANN counterpart of the best-performing SNN (MLP3_NonSpike, described in Section IV-C) was trained using the same dataset splits, preprocessing pipeline, and locked hyperparameters as the SNN. The comparison evaluates two aspects: (1) classification performance, measured by accuracy and macro-F1 on the held-out test set, and (2) estimated energy consumption per inference.

Energy was estimated using the 45 nm hardware model of Yao et al. [38], with $E_{\text{MAC}} = 4.6\,\text{pJ}$ per multiply-accumulate and $E_{\text{AC}} = 0.9\,\text{pJ}$ per accumulate operation. This comparison provides insight into the potential benefits of using spiking models for tactile perception on resource-constrained platforms.

Because both models receive the same z-score normalized (analog) inputs, the first-layer costs are identical. The energy comparison therefore focuses on the internal computations, summarized in Table IV. The key difference lies in fc2: the ANN performs dense MAC operations at every timestep, while the SNN receives binary spike inputs, reducing fc2 to sparse accumulate operations scaled by the hidden firing rate r_{LIF}. A second, hypothetical condition simulates an event-based sensor using delta encoding: a spike is emitted when the change in the z-scored signal exceeds a threshold $\theta_\delta = 0.1$, with separate channels for positive and negative crossings. The threshold was chosen to produce a sufficient spike rate for accurate classification while remaining small enough to capture meaningful signal changes, verified via a preliminary test. This converts the analog input to binary spike trains, so fc1 becomes a sparse AC operation rather than a MAC. Full details are provided in Appendix B-C.

VI. RESULTS

A. Datasets and Preprocessing

Representative processed pressure signals, preprocessing step visualizations, noise augmentation, and backward-sweep signals are provided in Appendix A.

B. Preprocessing Ablation

Table V summarizes the preprocessing ablation results over three random seeds.

TABLE V
PREPROCESSING ABLATION.

Condition	Val. acc. (%)
C1_N0	88.89 ± 0.0
C0_N0	33.33 ± 0.0
C1_N1	**92.59 ± 5.2**

C0/C1: uncropped/cropped to contact phase; N0/N1: min–max/z-score normalization. Evaluated using the MLP3 baseline with Adam, batch size 16, and 200 epochs with early stopping. The validation set contained 18 samples (one sample ≈ 5.6 pp). Cropping was essential (C0_N0 at chance); z-score normalization (C1_N1) outperformed min–max.

The results show that cropping is essential for successful classification. The uncropped condition, C0_N0, performs at chance level across all three seeds, with a validation accuracy of 33.33% and a macro-F1 of 16.67%. This indicates that, without restricting the sequence to the contact-relevant segment, the input is dominated by the approach phase, which contained little or no object-discriminative pressure information. Under this condition, the model converges to a trivial majority-class solution. Cropping is therefore retained as a fixed component of the preprocessing pipeline.

Among the cropped conditions, z-score normalization outperforms min–max normalization. Condition C1_N1 achieves the highest mean validation accuracy and macro-F1, reaching 92.59% and 92.38%, respectively, compared with 88.89% and 88.57% for C1_N0.

The variance observed for C1_N1 is higher than for C1_N0, although this should be interpreted with caution given the small validation set. Since one validation sample corresponds to 5.6 percentage points, the difference between 88.89% and 100% corresponds to only two additional correctly classified samples. The higher standard deviation therefore likely reflects small-sample variation rather than a qualitatively different performance regime.

C. Model Selection

Table VI compares the five fully spiking model families evaluated in the constrained hyperparameter sweep.

Among the five model families, MLP3 achieved the highest mean validation accuracy, reaching $97.22 \pm 7.1\%$, while also requiring by far the fewest parameters. In contrast to the convolutional and recurrent variants, MLP3 processed each time step as a flat sensor vector and relied on LIF membrane dynamics for temporal integration. This indicated that the instantaneous whisker pressure pattern already contained sufficient discriminative information, while more structured temporal connectivity provided limited additional benefit on this dataset.

A notable result was that the validation set became saturated for the strongest models. MLP3, MLP4, TCN, and TCN_Whisker all achieved at least one run with 100% validation accuracy, and within the MLP3 family a large fraction of runs reached this ceiling. Because the validation split contained only 18 samples, it had limited power to discriminate between top-performing configurations. For this reason, model selection within the best-performing group was based not only on peak validation score, but also on robustness and parameter efficiency.

RNN performed substantially worse than the other architectures, with a mean validation accuracy of $65.81 \pm 15.5\%$ and a maximum of only 88.89%, indicating that recurrent structure did not provide a systematic advantage in this setting.

For the convolutional architectures, temporal max-pooling consistently improved performance. Applying pooling with $k = 2$, which halved the temporal dimension before classification, yielded better results than the non-pooled variants for both TCN and TCN_Whisker.

From a deployment perspective, MLP3 was also the most parameter-efficient architecture. The smallest MLP3 configuration, using 901 parameters, achieved the same peak validation accuracy as substantially larger alternatives, including MLP4 and the convolutional models. Based on these results, MLP3 was selected as the baseline architecture for the subsequent experiments. Because multiple MLP3 configurations reached the validation ceiling, the final locked configuration was chosen using parameter count as a tiebreaker, resulting in the selection of the smallest reproducibly high-performing model. The training dynamics of the selected model are shown in Fig. 5.

D. Baseline Evaluation

To verify the robustness of the selected configuration, the locked MLP3 model was retrained using three random seeds. All three runs converged successfully and achieved 100% validation accuracy, confirming that the selected configuration was reproducible and not an artifact of a single favorable initialization.

Final evaluation on the held-out test set further supported this choice. The locked MLP3 configuration achieved 100.0 ± 0.0% test accuracy and 100.0 ± 0.0% macro-F1 across three seeds on the 24-sample test set. This confirmed that the selected baseline generalized perfectly under i.i.d. evaluation.

A well-tuned TCN also reached 100% test accuracy and macro-F1, indicating that convolutional spiking models could match MLP3 on the held-out test split. However, MLP3 remained the preferred baseline because it combined perfect test performance with substantially lower parameter count and more reliable training behavior.

E. In-depth Analysis of Selected Model

An example of the layer-wise spiking activity for a single cube trial is shown in Fig. 6. The input layer (top) displays the z-score normalized pressure signals as a heatmap across all 24 channels. The hidden layer (middle) shows the resulting spike raster for the 32 LIF neurons, illustrating the sparse, temporally distributed activity that emerged from the analog input. The output layer (bottom) shows the spikes of the three class neurons; the sustained firing of output neuron 0 indicated a correct classification as cube.

F. Classification Performance

1) Effect of Finite Difference Features: Using the selected MLP3 baseline, a controlled feature ablation was conducted to

TABLE VI
SPIKING ARCHITECTURE COMPARISON ACROSS TRAINING RUNS.

Spiking Architecture	Mean val. acc. (%)	Best val. acc. (%)	No. of params
MLP3	**97.22 ± 7.1**	**100.0**	**901**
MLP4	89.29 ± 13.9	100.0	20,099
TCN	87.72 ± 11.7	100.0	22,883
TCN_Whisker	87.65 ± 16.5	100.0	9,799
RNN	65.81 ± 15.5	88.89	14,287

150 random-search trials across five fully spiking model families with the fixed z-score preprocessing pipeline. Validation set contained 18 samples (one sample ≈ 5.6 pp). No. of params refers to the best-performing configuration per family. MLP3 achieved the highest mean accuracy with the fewest parameters and was selected as the baseline for all subsequent experiments.

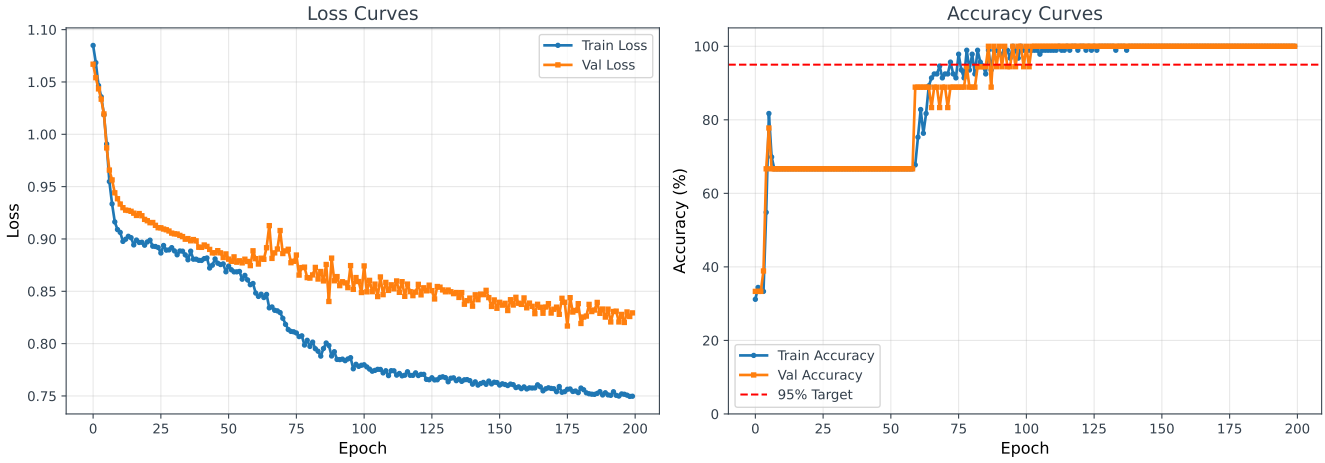

Fig. 5. Training and validation accuracy and loss curves for the selected MLP3 model. The model reached stable high accuracy from approximately epoch 75 onward, while the validation loss did not diverge from the training loss, indicating no clear evidence of overfitting.

TABLE VII
HELD-OUT TEST-SET PERFORMANCE OF THE TOP SELECTED RUNS.

Model	Test acc. (%)
MLP3	**100.0 ± 0.0**
MLP4	69.4 ± 2.0
TCN	100.0 ± 0.0

Test set contained 24 samples (one sample ≈ 4.2 pp). Each model was retrained with three seeds using the best configuration from the sweep. MLP3 and TCN both achieved perfect test performance; MLP4 overfitted despite matching MLP3 on validation.

TABLE VIII
EFFECT OF FINITE DIFFERENCE FEATURES ON CLASSIFICATION PERFORMANCE.

Condition	Description	Test acc. (%)
D0	Raw pressure only	**100.0 ± 0.0**
D1	Finite difference only	94.4 ± 5.2
D2	Raw + finite difference	**100.0 ± 0.0**

D0: 24 raw pressure channels; D1: 24 finite difference channels ($\Delta X[t] = X[t] - X[t-1]$, with $\Delta X[0] = 0$); D2: 48 channels (raw + finite difference concatenated). D0 and D1 used the locked MLP3 baseline (901 parameters); D2 used 1,669 parameters due to the doubled input size. Raw pressure alone was sufficient; finite difference features were informative but redundant when combined with raw input. All conditions evaluated on the same 24-trial held-out test set.

evaluate whether finite difference features improved classification performance. Only the input representation was varied, while the preprocessing pipeline, architecture, and training settings were kept fixed. Representative finite difference signals for the three object classes are shown in Appendix A, Fig. 14.

The results showed that raw pressure alone was sufficient to achieve ceiling performance on this dataset. Condition D0 reached 100.0% validation accuracy and 100.0% test accuracy across all three seeds, confirming that the selected baseline was stable and that the raw whisker pressure signals already contained enough information for perfect classification under the present evaluation setting.

Using only finite difference features led to slightly lower performance. Condition D1 achieved 94.4% validation accuracy and 94.4% test accuracy on average, indicating that the temporal rate of change of the pressure signal was informative, but not as informative as the raw pressure sequence itself. Since performance remained well above the three-class chance level of 33.3%, finite difference features clearly retained class-discriminative information.

Combining raw and finite difference features did not improve performance over the raw baseline. Condition D2 matched D0 exactly, reaching 100.0% on both validation and test sets. This indicated that, for the present task, the finite difference representation did not provide additional discriminative information beyond what was already available in the raw pressure sequence.

2) Temporal Window: Fig. 7 shows the performance of the selected MLP3 model under different temporal truncation

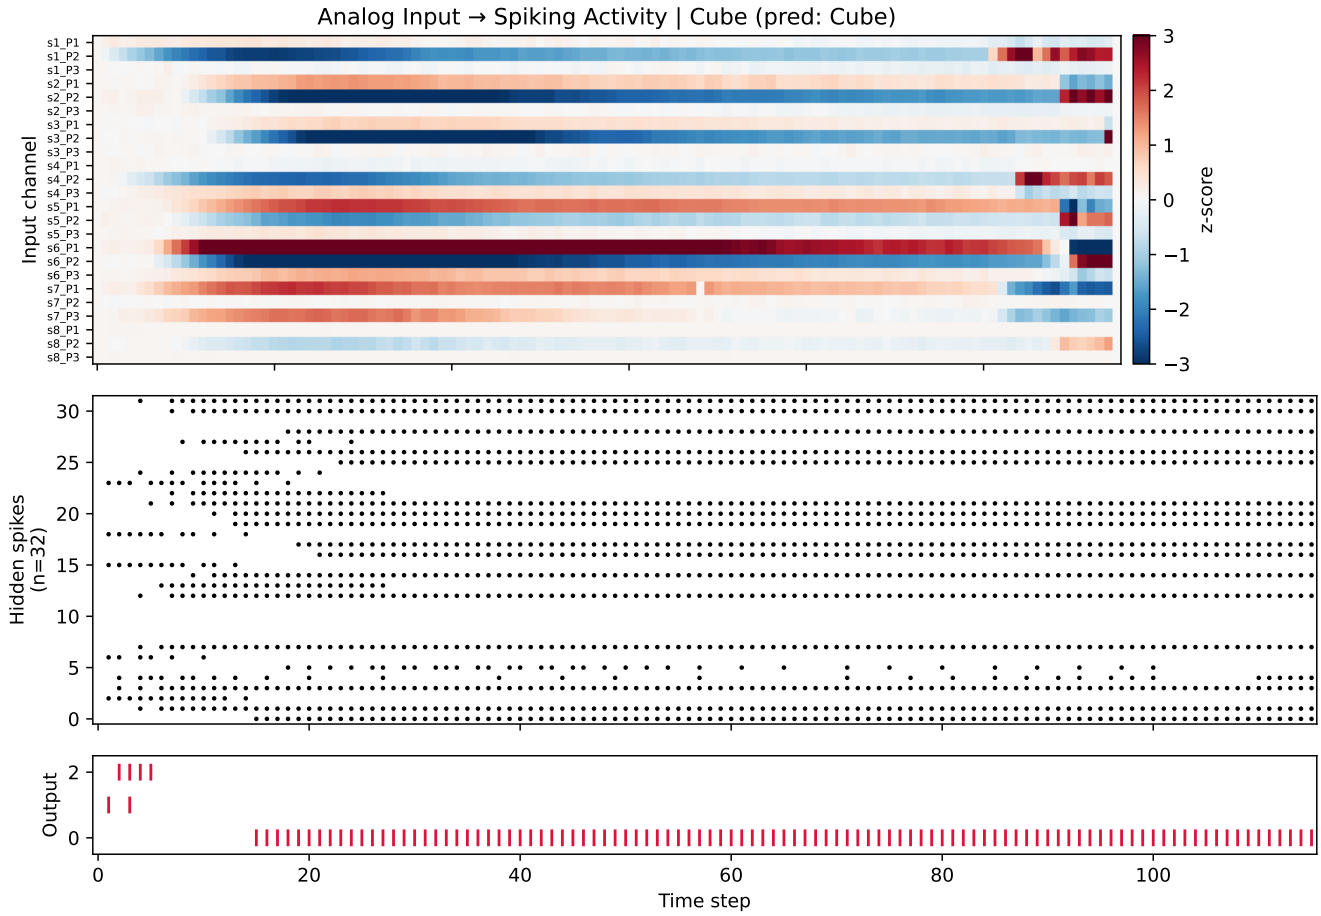

Fig. 6. Layer-wise activity of the selected MLP3 model for a single cube trial. Top: z-score normalized analog input across 24 pressure channels (8 whiskers × 3 channels); several channels remain near zero throughout, reflecting whiskers with low sensitivity or weak contact response. Middle: spike raster of the 32 hidden LIF neurons. Several hidden neurons (indices 8–11) are silent throughout the trial, indicating dead neurons that do not contribute to classification. The remaining hidden neurons fire at a steady rate after an initial transient of approximately 15 timesteps. Bottom: output layer spikes for the three class neurons (0: cube, 1: sphere, 2: rectangular prism). Neurons 1 and 2 fire briefly at the start before becoming silent, while neuron 0 fires continuously from around timestep 15 onward, showing that the model commits to the correct cube classification early in the sweep — well before the characteristic pressure release at the end.

conditions. Both head (first N timesteps) and tail (last N timesteps) segments of the sequence were evaluated to determine where discriminative information was concentrated. Each condition was trained from scratch on the S1-only subset (45 trials at fixed speed), with a 9-sample test set.

For the head conditions, a clear threshold was observed. Using only the first 10 timesteps (∼2 s of contact) resulted in reduced performance (70.4 ± 5.2% test accuracy), indicating that the very earliest signals were insufficient for reliable classification. However, performance improved sharply at 20 timesteps (∼4 s), where the model achieved 100% test accuracy across all three seeds. The head conditions remained at ceiling performance through 60 timesteps, confirming that the first 20 timesteps contained sufficient discriminative information.

The tail conditions demonstrated even stronger performance: the final 10 timesteps alone were sufficient to achieve 100% test accuracy, and performance remained at ceiling through 60 timesteps. This indicated that discriminative object information was also present during the late-contact phase, when the whiskers were fully engaged with the object.

Apparent degradation was observed at the largest windows (H80, H_full, and T80). However, these drops corresponded to only one or two misclassifications and should be interpreted with caution given the small test set: on 9 samples, a single error produced an 11.1 percentage point drop. Since a fresh model was trained for each condition, these variations were attributed to stochastic optimization on the small training set rather than to a systematic effect of window length. Notably, this degradation at H_full does not reflect an inherent limitation of longer sequences: the main model, trained on the full dataset (135 trials) with the full sequence length, achieved 100% test accuracy, confirming that full-length sequences are sufficient when sufficient training data is available.

3) Whisker Ablation: Table IX summarizes the performance of the selected MLP3 model under different input configurations, ranging from the full whisker array (S_all) to single-whisker inputs (S0 to S7). The physical positions of the whiskers in the array are shown in Fig. 3.

The full-array condition (S_all) achieved 100.0 ± 0.0% test

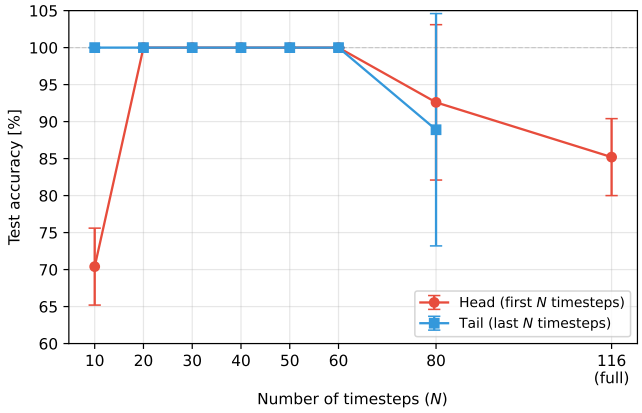

Fig. 7. Test accuracy as a function of the number of input timesteps N for head (first N) and tail (last N) truncations. Head conditions required at least 20 timesteps ($\sim$4 s) to reach ceiling performance; tail conditions achieved 100% accuracy with as few as 10 timesteps. The degradation at $N = 80$ and $N = 116$ (full) was attributed to small-sample variability (9-sample test set, one error $\approx$ 11.1 pp). All conditions used the locked MLP3 baseline on S1-only trials (45 trials, fixed speed). Error bars show $\pm$1 std over three seeds.

TABLE IX
WHISKER ABLATION RESULTS (RQ-E).

Condition	Channels	Test Acc (%)
S_all	24	**100.0 ± 0.0**
S0	3	79.2 ± 5.9
S1	3	86.1 ± 10.9
S2	3	75.0 ± 6.8
S3	3	**90.3 ± 2.0**
S4	3	84.7 ± 3.9
S5	3	88.9 ± 2.0
S6	3	62.5 ± 3.4
S7	3	43.1 ± 3.9

S_all: all 8 active whiskers (24 channels); S0–S7: individual whiskers (3 pressure channels each). MLP3 input layer adapted to 3 channels for single-whisker conditions; all other settings identical to the locked baseline. Test set contained 24 samples. Best single whiskers (S3, S5) reached $\sim$90%; no single whisker matched the full array.

accuracy across all seeds, confirming the reproducibility of the selected baseline. All single-whisker models converged successfully, indicating that the SNN was able to extract discriminative information even from a single whisker, albeit with reduced performance and higher variability.

Performance varied substantially across individual whiskers. The best-performing single whiskers (S3 and S5) achieved approximately 90% test accuracy, approaching the performance of the full array. In contrast, weaker whiskers such as S6 and S7 achieved significantly lower performance (approximately 63% and 43%, respectively), indicating that not all sensors contributed equally to the classification task.

Despite the strong performance of the best individual whiskers, none matched the full-array performance. The remaining gap suggested that combining multiple whiskers provided complementary spatial information that could not be fully recovered from any single sensor.

G. Generalization

1) Depth and Speed: The domain generalization experiments evaluated whether the selected MLP3 model transferred to unseen indentation depths and sweep speeds. The results are shown in Fig. 8. Speed generalization was nearly perfect even without augmentation. In all leave-one-speed-out (LOSO) folds, the model trained on two speeds generalized to the held-out speed with at least 99.3% test accuracy. This indicated that variation in sweep speed did not constitute a meaningful domain shift for the whisker pressure signals in this task.

Depth generalization was more challenging. In the leave-one-depth-out (LODO) setting, performance varied substantially across folds. The shallow depth condition (D1) was the most difficult, achieving 77.1% test accuracy without augmentation, while D3 reached 88.9%. The intermediate depth (D2) generalized perfectly.

Data augmentation consistently improved or matched the baseline performance. Across all folds, additive noise augmentation (A1) performed at least as well as the non-augmented condition (A0). The largest improvement occurred for the hardest fold (LODO D1), where augmentation increased accuracy from 77.1% to 85.2%. In folds where the baseline already reached ceiling performance, augmentation had little effect.

Overall, these results indicated that sweep speed was not a meaningful source of domain shift, while indentation depth –particularly shallow contact– remained the primary generalization challenge. Noise augmentation reduced this gap without degrading performance in easier domains. The confusion matrices in Appendix B-A, Fig. 17, reveal that the primary failure mode was cube–corner confusion: at the shallow depth (D_1), cube was frequently misclassified as corner, while at the deep depth (D_3), the pattern reversed. Sphere was never confused regardless of depth.

2) Direction: The direction generalization experiment evaluated whether the selected MLP3 model trained on forward sweeps transferred to backward sweeps. The results are summarized in Table X; full numerical results including macro-F1 are provided in Table XII in Appendix B-B. Note that all results in this section were obtained using the locked MLP3 baseline; other architectures may exhibit different transfer behavior.

The backward sweep dataset was intrinsically classifiable. In condition C1, where training and testing were both performed on backward data, the model achieved 100.0% test accuracy across all three seeds. This showed that the backward dataset was clean and that the classification task remained fully solvable under reverse motion.

In contrast, zero-shot transfer from forward to backward motion failed completely. In condition C2, the model was trained only on forward sweeps and tested on backward sweeps, yielding a mean test accuracy of 24.4% and a mean macro-F1 of 12.9%, both well below the backward-data baseline. This indicated that the features learned by the MLP3 from forward sweeps did not directly transfer to reverse motion.

Joint training substantially restored performance. In condition C3, where the model was trained on both forward and backward sweeps, test accuracy increased to 88.9% and test macro-F1 to 85.2% on average. This recovery showed

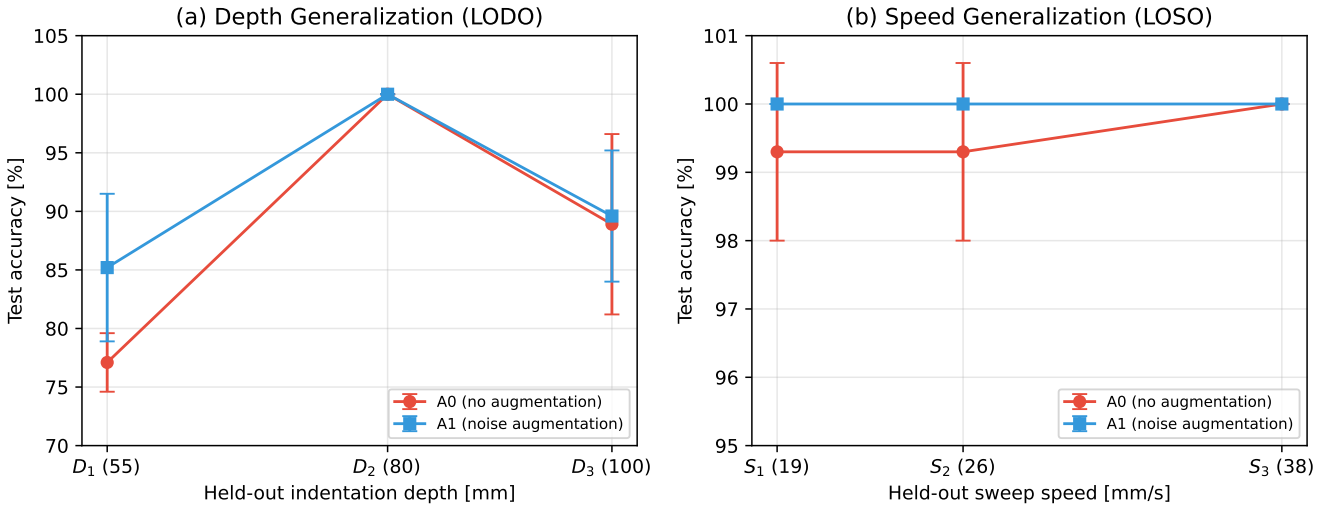

Fig. 8. Domain generalization accuracy across indentation depth (a) and sweep speed (b). LODO: leave-one-depth-out; LOSO: leave-one-speed-out. A0: no augmentation; A1: additive Gaussian noise jitter ($\sigma = 0.0064$). All models used the locked MLP3 baseline trained from scratch per fold. Error bars show ± 1 std over three seeds. Speed generalization was near-perfect; shallow depth (D_1) was the hardest domain, improved by augmentation.

TABLE X
DIRECTION GENERALIZATION RESULTS.

Condition	Source val. acc. (%)	Test acc. (%)
C1 (backward only)	100.0 ± 0.0	100.0 ± 0.0
C2 (forward→backward)	100.0 ± 0.0	24.4 ± 6.8
C3 (joint training)	88.9 ± 15.7	88.9 ± 15.7

C1: trained and tested on backward data only; C2: trained on forward data, tested zero-shot on backward data; C3: trained jointly on forward and backward data, tested on backward. All conditions used the locked MLP3 baseline. Source val. acc. reflects performance on the source-domain validation split (forward for C2, backward for C1/C3) and is reported to confirm successful training, not as a measure of target-domain generalization. Backward test set contained 6 samples (one sample ≈ 16.7 pp). Forward→backward zero-shot transfer failed; joint training recovered most performance.

that the direction shift could be mitigated when both motion directions were represented during training. The remaining variance across seeds was amplified by the small backward test set (6 samples), where a single misclassification produced a large change in percentage score.

H. Comparison with ANN Model

The SNN MLP3 was compared against the structurally equivalent non-spiking ANN (MLP3_NonSpike, described in Section IV-C). Both models were trained with the locked hyperparameters on the canonical split over three seeds.

Classification performance. Both models achieved 100% test accuracy (mean ± std over three seeds), confirming that the ANN baseline was equally capable on this task. The SNN therefore matched ANN performance while operating through spiking dynamics.

Estimated energy consumption. Because the SNN received analog (z-scored) inputs, the first layer still performed multiply-accumulate (MAC) operations identical to the ANN; the energy advantage of accumulate-only (AC) operations applied only from the hidden layer onward.

Under this analog-input scenario, the mean SNN total energy was 314.9 nJ per inference, compared to 347.8 nJ for the ANN — a reduction of approximately 9.5%. The saving arose from the output layer: the SNN performed AC operations driven by sparse hidden spikes (mean firing rate $r_{\text{LIF}} = 0.39 \pm 0.08$), consuming 3.0 nJ, while the ANN output layer required 38.6 nJ in MACs. Fig. 9 illustrates this fc2 energy difference.

The delta-encoded condition (see Section IV-C) reduced total energy to 17.6 nJ, a reduction of approximately 95% compared to the ANN. This represents an upper bound on SNN energy efficiency for this architecture if spike-encoded sensors were used. The full numerical comparison is provided in Table XIII in Appendix B-C, and Fig. 18 shows the per-layer energy breakdown for all three scenarios.

VII. DISCUSSION AND LIMITATIONS

This study presented the first application of spiking neural networks to tactile pressure data from a multi-whisker sensor array, addressing a gap in neuromorphic tactile sensing where prior work had focused exclusively on fingertip-style sensors. The results showed that SNN-based classification was feasible under controlled conditions. The selected model generalized well across sweep speed, less well across indentation depth, and poorly across sweep direction unless both directions were represented during training. This indicated that the learned representation was robust to some nuisance factors, but remained sensitive to motion conditions that substantially altered the temporal structure of the whisker signals. The direction sensitivity is expected for an MLP: because a fully connected layer assigns a fixed weight to each input channel, the model learns a spatial map associating specific sensor positions with specific objects. Reversing the sweep direction changes which whiskers make first contact and inverts the temporal ordering of pressure signals across channels, so the learned weights no longer correspond to the correct spatial relationships. Joint

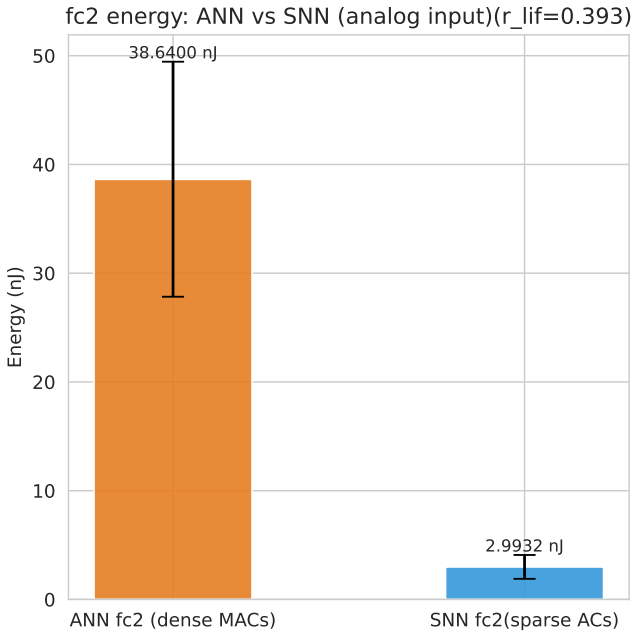

Fig. 9. Output layer (fc2) energy comparison between the ANN and SNN under analog input. The ANN performed dense MAC operations at every timestep, while the SNN fc2 layer received sparse binary spikes (mean hidden firing rate $r_{\mathrm{LIF}} = 0.393$), reducing fc2 energy by approximately $13\times$.

training on both directions resolves this by exposing the model to both spatial mappings.

A notable finding was that the simplest architecture a three-layer fully connected SNN with only 901 parameters outperformed more complex convolutional and recurrent alternatives. Although the MLP3 lacked explicit temporal connectivity, its high membrane decay factor ($\beta = 0.938$) enabled the LIF neurons to retain information across many timesteps through their membrane dynamics, effectively providing implicit temporal integration. The results suggested that this implicit temporal memory, combined with the spatial pressure pattern across whiskers, was sufficient for coarse object classification, and that the additional structural complexity of convolutional or recurrent layers provided limited benefit on this dataset. The temporal window experiment supported this interpretation: as few as 20 timesteps from the onset of contact were sufficient for perfect classification, indicating that the discriminative signal appeared early and did not require integration over the full sweep duration. Notably, discriminative information was present in both early and late contact phases: as few as 10 timesteps from the end of the sweep were equally sufficient for perfect classification, suggesting that the release phase carries geometry-specific signatures that are as informative as the initial contact transient. This indicates that object identity is encoded across the full sweep, not only at contact onset. The whisker ablation further showed that classification was possible from a single whisker, although the full array provided the most reliable performance. The depth generalization experiments revealed that the primary failure mode was confusion between cube and corner at extreme indentation depths, likely because both objects share a flat surface and edge geometry, which the model finds difficult to distinguish under varying contact conditions. The preprocessing results also offered an insight specific to analog SNN inputs: z-score normalization outperformed min-max scaling, which may be explained by the fact that min-max compression restricts all values to $[0, 1]$, discarding the relative magnitude differences that carry information about increasing and decreasing pressure patterns, whereas z-score normalization preserves these differences and centers the distribution around zero, naturally producing both excitatory and inhibitory drive and improving gradient flow during training.

The energy comparison showed that the SNN achieved a modest 9.5% reduction in estimated energy relative to a structurally equivalent ANN when both received analog input. This limited saving arose because the first layer which dominated the energy budget performed identical MAC operations in both models. The advantage of spike-driven accumulate operations applied only to the output layer, where the hidden firing rate of $r_{\mathrm{LIF}} \approx 0.39$ reduced fc2 energy by approximately $13\times$. A hypothetical delta-encoded input would reduce total energy to 5.1% of the ANN baseline by converting the first layer to sparse AC operations. However, this energy estimation followed a simplified theoretical model [38] that ignores hardware implementation details and the temporal dynamics of spiking neurons; actual energy savings on neuromorphic hardware may differ. Nevertheless, this type of simplified estimation is considered useful for algorithm-level comparison and design guidance [34].

The broader neuromorphic tactile literature provides useful context. Prior SNN-based tactile studies have demonstrated high classification accuracy and substantial energy reductions on dedicated neuromorphic hardware, for example in texture recognition [23] and Braille character reading [24]. The present results are broadly consistent with those findings in showing that SNNs can match conventional network accuracy on tactile time-series tasks. Unlike those studies, the present work includes a direct comparison with a structurally equivalent non-spiking baseline, confirming that the spiking model matched ANN classification accuracy on the same task and dataset while offering an estimated reduction in inference energy. However, a direct quantitative comparison with prior neuromorphic tactile work is not meaningful, as those studies used different sensor modalities, task structures, and hardware platforms. The contribution of the present work lies specifically in extending spiking tactile classification to the spatially distributed, sweep-based signals of a multi-whisker array, a sensing paradigm not previously evaluated with SNNs.

Several limitations of the sensing setup should be noted. First, the whisker array exhibited channel-level variability, and two whisker units had to be excluded because of non-functional or low-sensitivity measurements. More generally, differences in sensor sensitivity across whiskers may have reduced the effectiveness of global normalization, since channels with smaller dynamic range could be compressed relative to more responsive channels. This also offers a likely explanation for the variation in single-whisker performance observed in the whisker ablation experiment: whiskers with lower intrinsic sensitivity produced weaker signals regardless of object geom-

etry. Additionally, the vertical position of a whisker in the array affected contact quality with curved objects — for the sphere, center whiskers contacted the widest equatorial region and followed a smooth, consistent arc, whereas whiskers toward the top or bottom of the array may have slid off the curved surface, producing noisier and less discriminative signals. Second, the dataset was relatively small, particularly for the backward-sweep experiments, which limited statistical power and increased the variance of held-out evaluations.

A further limitation was that the current system was not fully neuromorphic end-to-end. Although the classifier was implemented as an SNN, the whisker array produced conventionally sampled pressure measurements rather than event-based output, and all experiments were executed offline on standard computing hardware. The approach therefore demonstrated the value of spiking processing for tactile sequences, but did not yet realize the full sensing and energy-efficiency benefits associated with neuromorphic hardware.

The conclusions were also bounded by the experimental regime. Data were collected under slow, quasi-static conditions and at relatively low sampling frequency, which appeared sufficient for coarse object classification but would be unsuitable for tasks that require higher-frequency information such as texture recognition. Moreover, the experiments were conducted under controlled laboratory conditions and did not account for real-world effects such as actuator-induced vibration, aerodynamic disturbances, gravity-dependent motion, repeated probing behavior, or whisker buckling during arbitrary contacts. Finally, the backward-sweep dataset covered only one depth-speed configuration, which limited how broadly the direction-generalization results could be interpreted.

VIII. Conclusions and Future Work

This work demonstrated that spiking neural networks can reliably classify objects from biomimetic whisker sensor data, with a compact three-layer spiking MLP achieving 100% test accuracy while matching a conventional ANN in performance and offering a path to substantially lower energy consumption. The following summarises the findings across three objectives: classification performance, generalization, and computational efficiency.

Regarding *classification performance*, a compact three-layer spiking MLP with only 901 parameters achieved 100% test accuracy on a three-class object classification task, outperforming more complex convolutional and recurrent SNN architectures. Raw pressure signals were sufficient for classification; temporal derivative features were informative but redundant. The temporal window experiment showed that as few as 20 timesteps ($\sim$4 s) from the onset of contact were sufficient for reliable classification, and the whisker ablation demonstrated that single whiskers could classify above chance, although the full array provided the best performance.

Regarding *generalization*, sweep speed had little effect on classification, while indentation depth introduced a moderate shift particularly at shallow depths, where cube and corner objects became confusable. Sweep direction represented the strongest domain boundary: zero-shot transfer from forward to backward motion failed, but joint training on both directions recovered most performance. These results indicated that generalization depended strongly on the motion regime under which data were collected.

Regarding *computational efficiency*, the SNN achieved a 9.5% reduction in estimated energy consumption compared to a structurally equivalent ANN under analog input, with the saving concentrated in the output layer due to sparse spiking activity. A hypothetical delta-encoded input scenario would reduce energy by approximately 95% compared to the ANN. The SNN therefore matched ANN classification performance while offering a path toward substantially lower energy consumption.

These findings establish that fully spiking architectures are a viable candidate for low-power tactile perception on resource-constrained robotic platforms, where the combination of biological inspiration, temporal processing, and energy efficiency offers a compelling alternative to conventional neural networks. Overall, the present system should be interpreted as a proof of concept for spiking tactile perception rather than as a complete demonstration of deployable end-to-end neuromorphic sensing.

Several directions for future work follow directly from this study. A first priority is validation on neuromorphic hardware such as Intel Loihi, including explicit measurement of inference latency and energy consumption under realistic deployment conditions. A second direction is the integration of event-driven whisker sensors, which would convert the first layer from dense MAC to sparse AC operations and could substantially reduce energy consumption beyond the analog-input scenario studied here. A third direction is to increase the sampling frequency beyond the current 5 Hz to enable higher-frequency tactile tasks such as texture recognition. Finally, extending the task toward surface mapping or 3D localization and deploying the system on a mobile robotic platform such as a UAV would allow evaluation under realistic disturbances including actuator-induced vibration, aerodynamic effects, and unconstrained contact dynamics.

APPENDIX A
TACTILE DATASETS AND SIGNAL PROCESSING

A. Whisker Sweep Signals

(a) Cube (b) Sphere (c) Rectangular prism

Fig. 10. Processed whisker pressure signals from the forward sweep dataset for the three object classes. (a) Cube: pressure gradually increases after contact and remains relatively stable before a sharp release when the whisker disengages. (b) Sphere: more varied pressure responses across whiskers due to the curved surface; central whiskers maintain sustained contact while outer whiskers exhibit more fluctuating signals. (c) Rectangular prism: pressure increases steadily as the whisker slides along the inclined surface before releasing at the edge.

B. Cropping Procedure

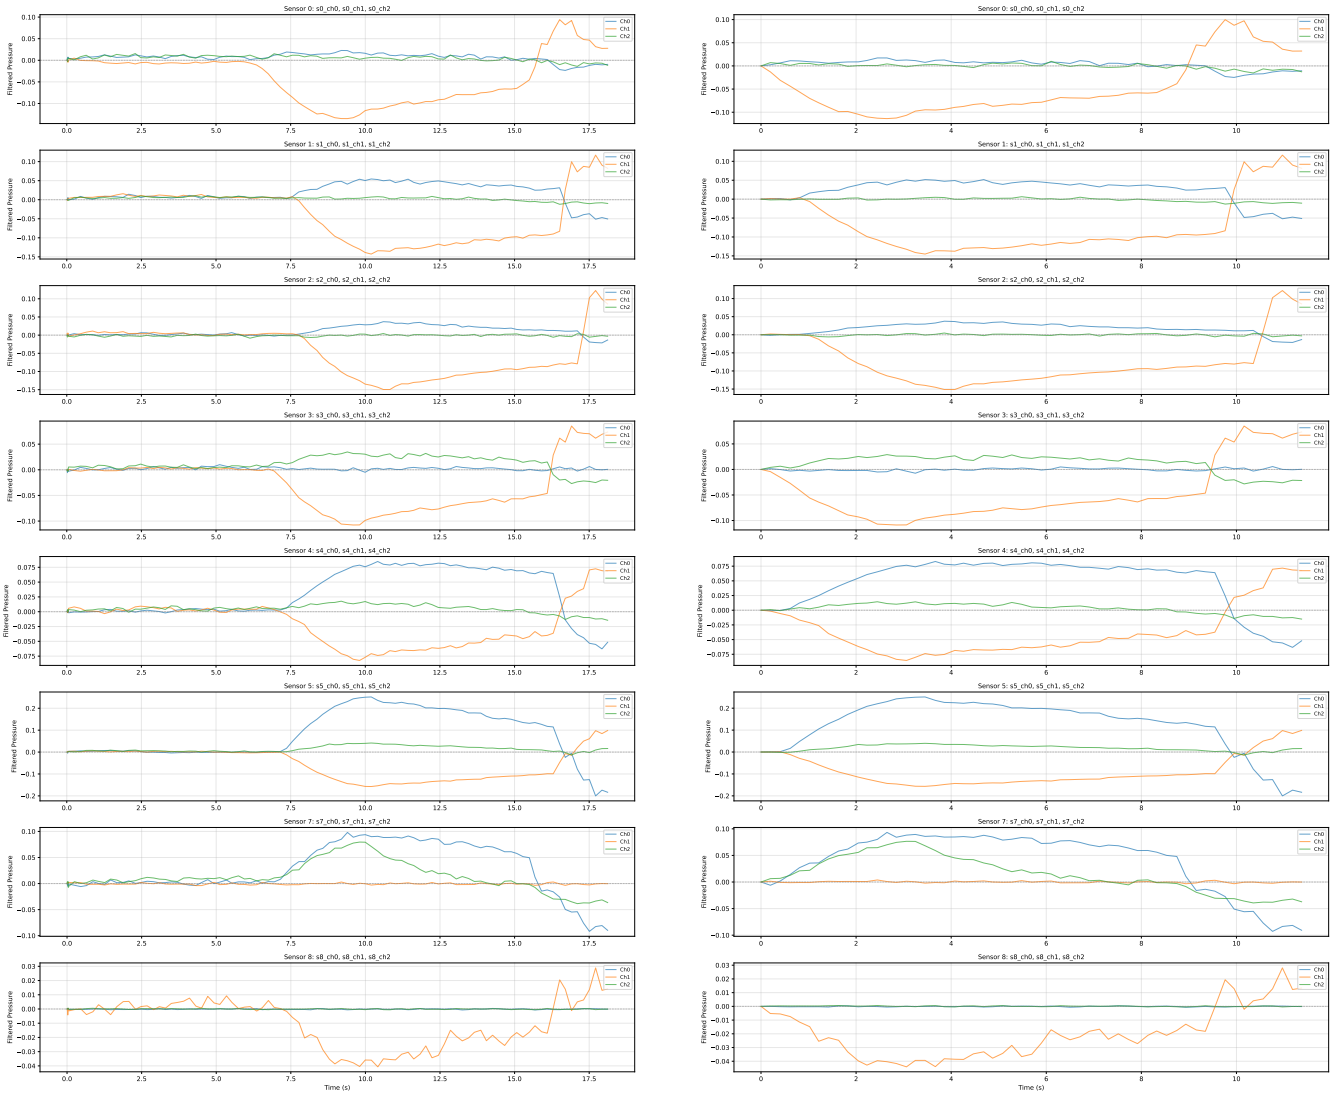

(a) Before cropping (full sequence)

(b) After cropping (contact phase only)

Fig. 11. Example of the cropping procedure applied to the whisker pressure signals. Using the recorded position data, the initial 250 mm of motion before contact is removed. This eliminates the phase where the sensor moves through air and retains only the contact-relevant portion of the signal used for training the network.

C. Before Baseline Removal

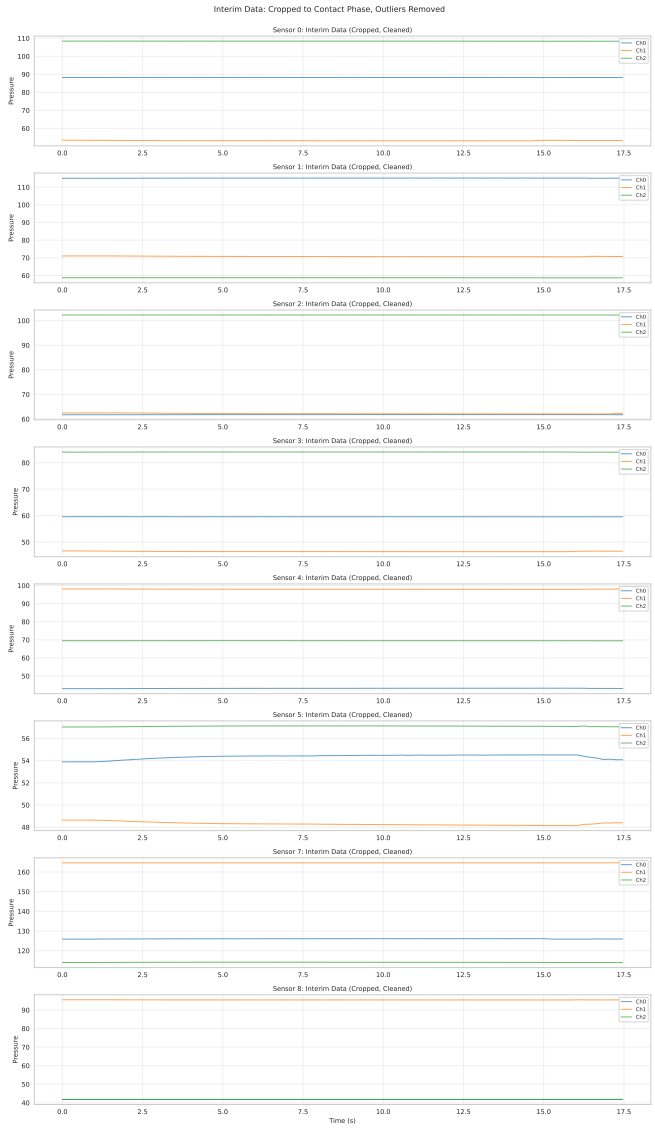

Fig. 12. Whisker pressure signals before baseline correction. Differences in baseline pressure between channels are substantially larger than the variations within each channel. Baseline correction was applied by subtracting the first valid value of each channel.

D. Band-pass Filtering

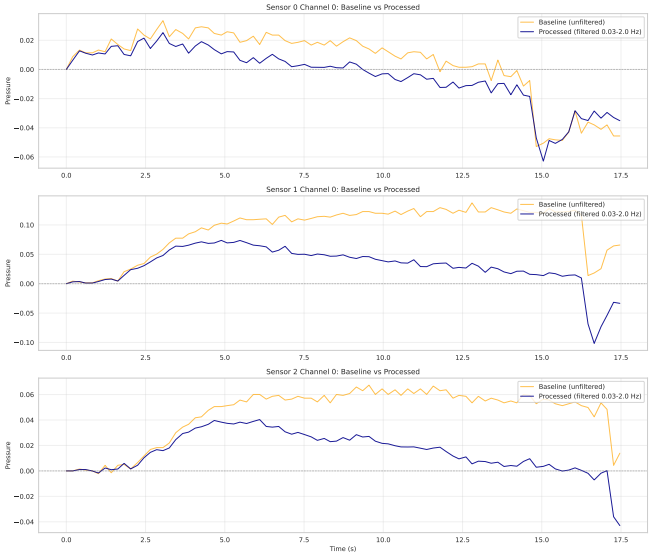

Fig. 13. Effect of the band-pass filter on three example channels. The high-pass component removed slow baseline drift, while the low-pass component attenuated high-frequency measurement noise.

E. Finite Difference Signals

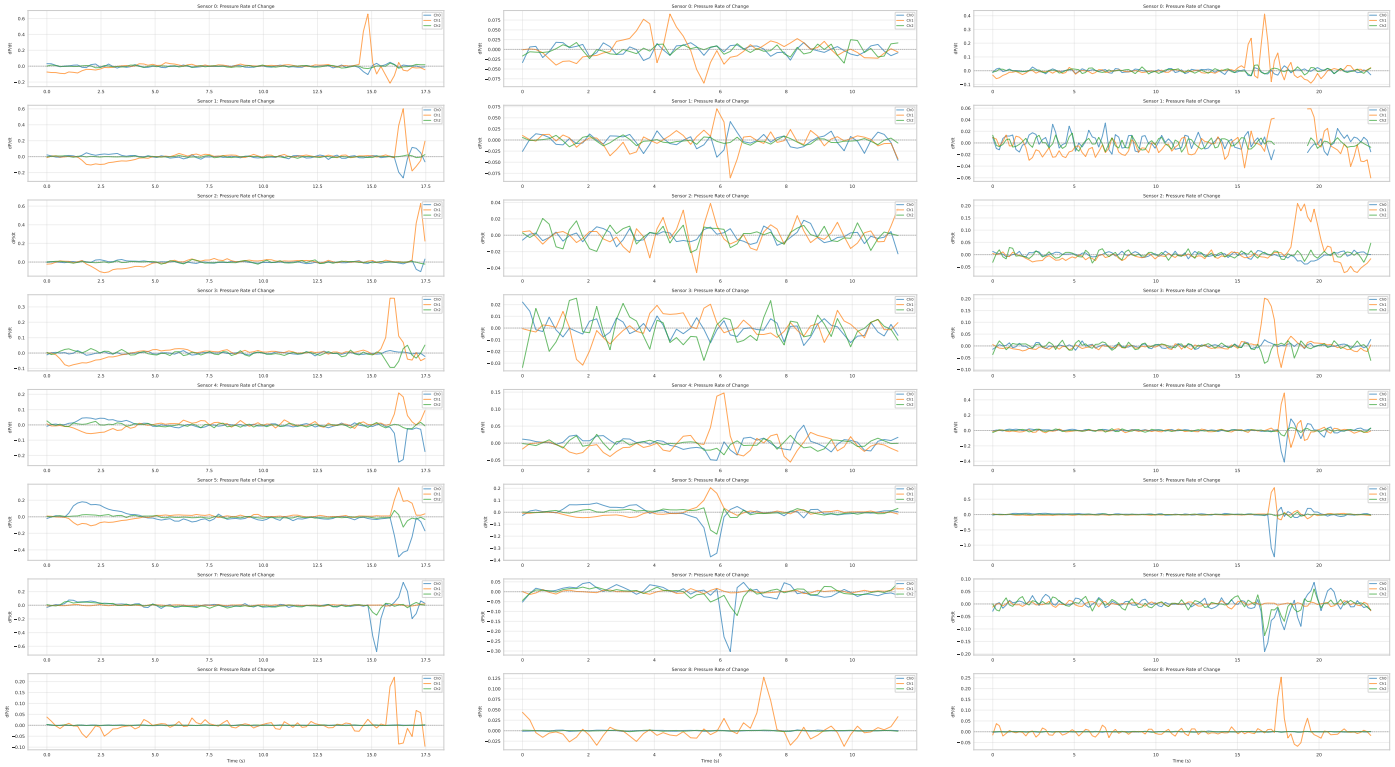

(a) Cube (b) Sphere (c) Rectangular prism

Fig. 14. First-order finite difference ($\Delta X[t] = X[t] - X[t-1]$) of whisker pressure signals from the forward sweep dataset for the three object classes. (a) Cube: gradual increase in the rate of change during contact and a large spike when the whisker suddenly releases from the object. (b) Sphere: more oscillatory finite difference signals due to the curved surface interaction, with the largest peaks occurring when the whiskers bend most strongly. (c) Rectangular prism: finite difference values remain small during the initial contact phase but increase sharply when the whiskers reach and slide off the edge.

F. Data Augmentation

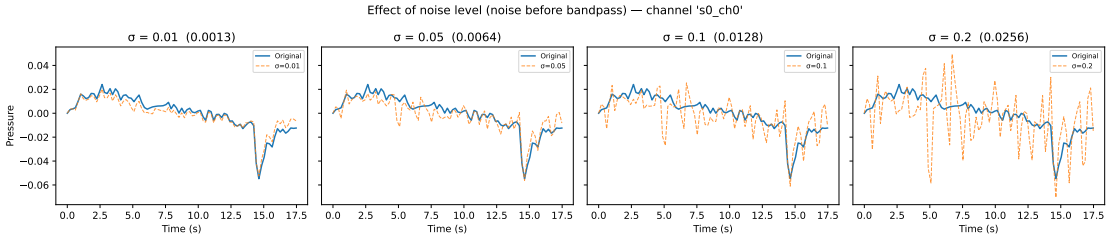

(a) Augmentation levels

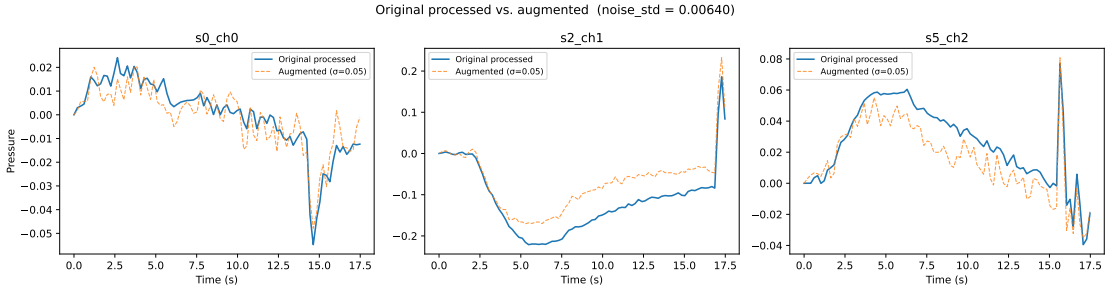

(b) Effect across channels

Fig. 15. Noise augmentation used in the depth and speed generalization experiments. (a) Gaussian noise with different standard deviations (σ) applied to the pressure signals. The selected value $\sigma = 0.0064$ provided sufficient variation while preserving signal structure. (b) Effect of the selected augmentation level across different pressure channels.

G. Backward Sweep Signals

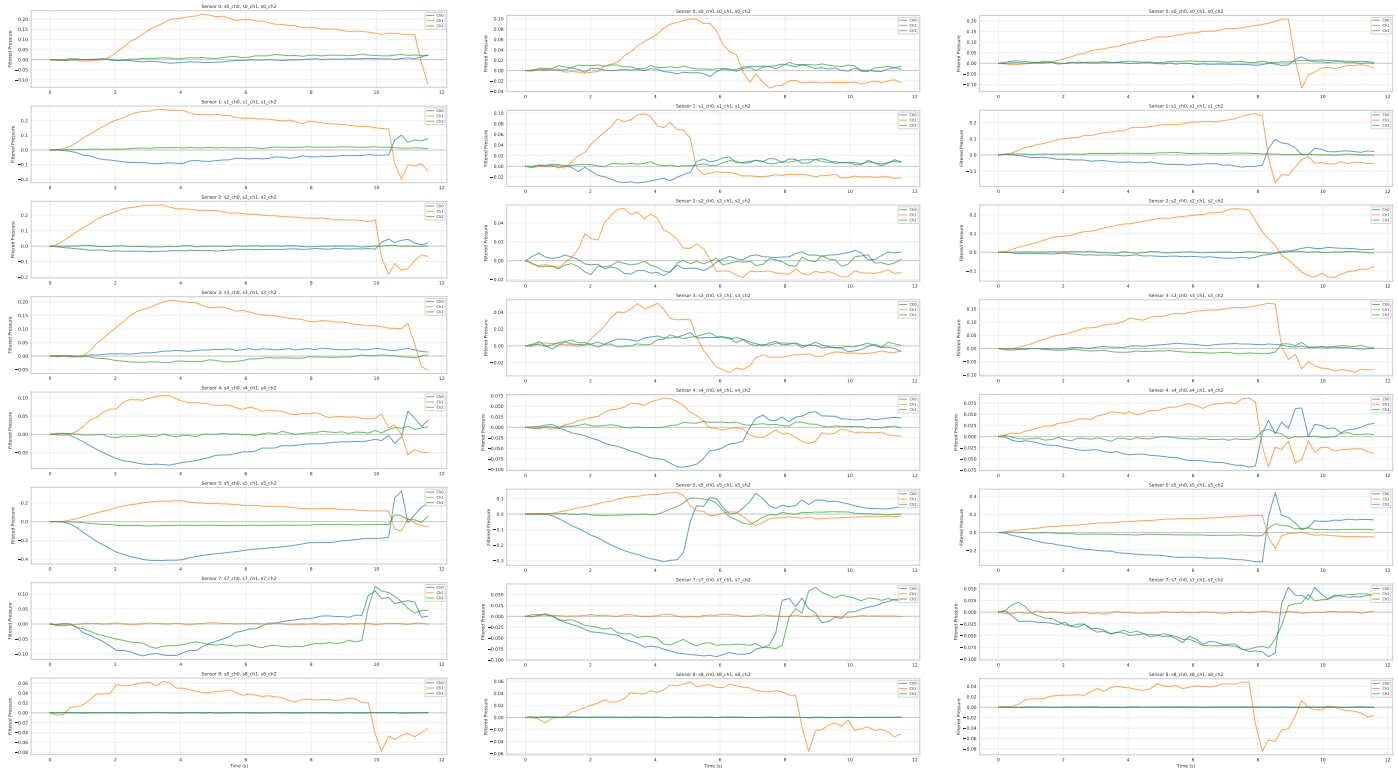

(a) Cube (b) Sphere (c) Rectangular prism

Fig. 16. Processed whisker pressure signals from the backward sweep dataset for the three object classes. These trials were collected at fixed depth D_2 and speed S_3 with the whisker array moving away from the object. Compared to the forward sweep signals (Fig. 10), the temporal structure is reversed: pressure release occurs at the start of the contact phase rather than at the end.

APPENDIX B
ADDITIONAL EXPERIMENTAL RESULTS

A. Depth and Speed Generalization

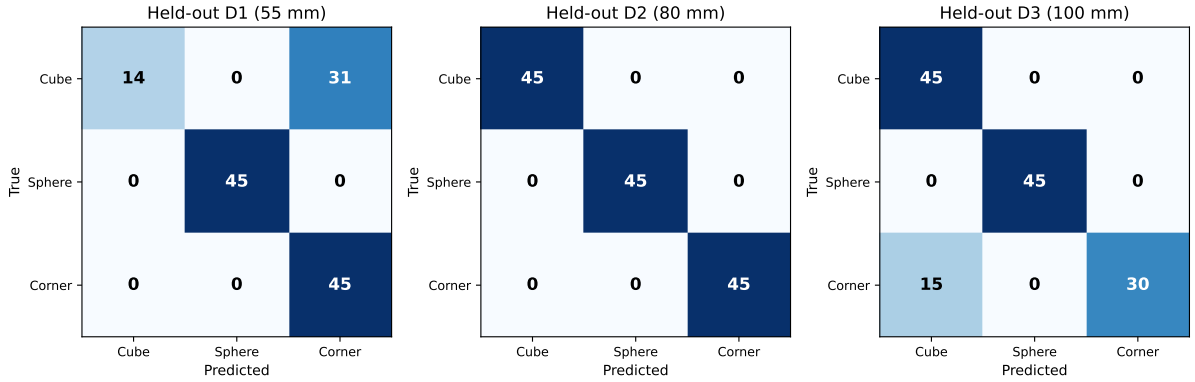

Fig. 17. Confusion matrices for the LODO depth generalization experiment (A0, no augmentation), aggregated across three seeds (45 samples per class). At shallow depth (D_1), cube was predominantly misclassified as corner. At deep depth (D_3), corner was confused with cube. At intermediate depth (D_2), all classes were perfectly separated. Sphere was never confused regardless of depth.

TABLE XI
DEPTH AND SPEED DOMAIN GENERALIZATION RESULTS ACROSS THREE RANDOM SEEDS.

Fold	Domain type	A0 Test Acc. (%)	A0 Test F1 (%)	A1 Test Acc. (%)	A1 Test F1 (%)
LODO testD1	Held-out depth D_1	77.1 ± 2.5	73.8 ± 3.6	85.2 ± 6.3	84.0 ± 7.5
LODO testD2	Held-out depth D_2	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
LODO testD3	Held-out depth D_3	88.9 ± 7.7	88.2 ± 8.3	89.6 ± 5.6	89.2 ± 6.0
LOSO testS1	Held-out speed S_1	99.3 ± 1.3	99.3 ± 1.3	100.0 ± 0.0	100.0 ± 0.0
LOSO testS2	Held-out speed S_2	99.3 ± 1.3	99.3 ± 1.3	100.0 ± 0.0	100.0 ± 0.0
LOSO testS3	Held-out speed S_3	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0

LODO: leave-one-depth-out; LOSO: leave-one-speed-out. A0: no augmentation; A1: additive Gaussian noise ($\sigma = 0.0064$). All models used the locked MLP3 baseline trained from scratch per fold. Results are mean ± std over three seeds.

B. Direction Generalization

TABLE XII
DIRECTION GENERALIZATION RESULTS ACROSS THREE RANDOM SEEDS.

Condition	Val. Acc. (%)	Test Acc. (%)	Test F1 (%)
C1 (backward only)	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
C2 (forward→backward)	100.0 ± 0.0	24.4 ± 6.8	12.9 ± 2.9
C3 (joint training)	88.9 ± 15.7	88.9 ± 15.7	85.2 ± 21.0

C1: trained and tested on backward data only; C2: trained on forward, tested zero-shot on backward; C3: joint training on both directions, tested on backward. All conditions used the locked MLP3 baseline. Backward test set contained 6 samples per seed (one misclassification ≈ 16.7 pp).

C. SNN vs. ANN Energy Comparison

TABLE XIII
SNN VS. ANN COMPARISON ON THE 24-SAMPLE TEST SET.

Metric	ANN	SNN (analog)	SNN (delta-enc.†)
Test accuracy	100 %	100 %	100 %
Input spike rate (r_δ)	—	—	0.203 ± 0.039
Activation density	0.508 ± 0.028	0.393 ± 0.075	0.393 ± 0.075
fc1 energy (nJ)	309.1	309.1	11.8
LIF energy (nJ)	—	2.8	2.8
fc2 energy (nJ)	38.6	3.0	3.0
Total energy (nJ)	347.8	314.9	17.6
Energy ratio vs. ANN	100 %	90.5 %	5.1 %

Energy estimated using 45 nm hardware model ($E_{\text{MAC}} = 4.6$ pJ, $E_{\text{AC}} = 0.9$ pJ) from Yao et al. Values are mean over 24 test samples. Activation density: fraction of active hidden neurons (ReLU for ANN, spike rate for SNN). r_δ: input spike rate under delta encoding ($\theta_\delta = 0.1$). †Hypothetical scenario: analog input replaced by spike-encoded signal using a change-detection threshold, converting fc1 from MACs to ACs. The analog SNN saves 9.5% total energy; delta encoding would reduce it to 5.1% of the ANN.

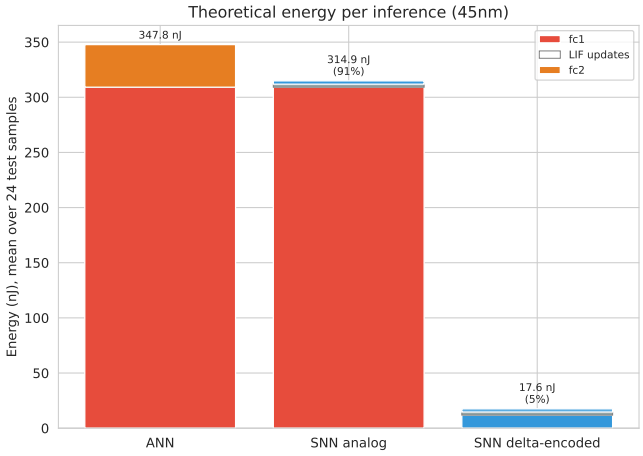

Fig. 18. Theoretical energy per inference (45 nm hardware model) for the ANN, SNN with analog input, and the hypothetical delta-encoded SNN. The stacked bars show fc1, LIF update, and fc2 contributions. The analog SNN saves 9.5 % via sparse fc2 operations; delta encoding reduces total energy by approximately 95%.

REFERENCES

[1] A. S. Ahl, "The role of vibrissae in behavior: A status review," *Veterinary Research Communications*, vol. 10, no. 1, pp. 245–268, Dec. 1, 1986, ISSN: 1573-7446. DOI: 10.1007/BF02213989. [Online]. Available: https://doi.org/10.1007/BF02213989 (visited on 05/20/2025).

[2] T. J. Prescott, M. J. Pearson, B. Mitchinson, J. W. Sullivan, and A. G. Pipe, "Whisking with robots," *IEEE Robotics & Automation Magazine*, vol. 16, no. 3, pp. 42–50, Sep. 2009, ISSN: 1558-223X. DOI: 10.1109/MRA.2009.933624. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/5233415 (visited on 06/02/2025).

[3] R. Russell, "Using tactile whiskers to measure surface contours," in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, May 1992, 1295–1299 vol.2. DOI: 10.1109/ROBOT.1992.220070. [Online]. Available: https://ieeexplore.ieee.org/document/220070 (visited on 06/19/2025).

[4] Active Touch Laboratory, *Active touch laboratory — whisker sensing*, YouTube. [Online]. Available: https://www.youtube.com/watch?v=bRZKczxk--4 (visited on 04/13/2026).

[5] A. Bredenbeck, C. D. Santina, and S. Hamaza, "Embodying compliant touch on drones for aerial tactile navigation," *IEEE Robotics and Automation Letters*, vol. 10, no. 2, pp. 1209–1216, Feb. 2025, ISSN: 2377-3766. DOI: 10.1109/LRA.2024.3519888. [Online]. Available: https://ieeexplore-ieee-org.tudelft.idm.oclc.org/document/10806756 (visited on 05/19/2025).

[6] M. A. Lin, E. Reyes, J. Bohg, and M. R. Cutkosky, "Whisker-inspired tactile sensing for contact localization on robot manipulators," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2022, pp. 7817–7824. DOI: 10.1109/IROS47612.2022.9982122. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9982122 (visited on 05/20/2025).

[7] J. H. Solomon and M. J. Z. Hartmann, "Extracting object contours with the sweep of a robotic whisker using torque information," *The International Journal of Robotics Research*, vol. 29, no. 9, pp. 1233–1245, Aug. 1, 2010, ISSN: 0278-3649. DOI: 10.1177/0278364908104468. [Online]. Available: https://doi.org/10.1177/0278364908104468 (visited on 06/02/2025).

[8] G. Kshettry, T. Patel, K. Tiwari, and A. Bera, "TuneNav: Tunable whisker array for touch-based navigation in confined spaces," in *2024 IEEE MIT Undergraduate Research Technology Conference (URTC)*, Oct. 2024, pp. 1–5. DOI: 10.1109/URTC65039.2024.10937626. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10937626 (visited on 06/30/2025).

[9] T. Kossas, W. Remmas, R. Gkliva, A. Ristolainen, and M. Kruusmaa, "Whisker-based tactile navigation algorithm for underground robots," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 13 164–13 170. DOI: 10.1109/ICRA57147.2024.10610762. [Online]. Available: https://ieeexplore.ieee.org/document/10610762/ (visited on 06/30/2025).

[10] M. Salman and M. J. Pearson, "Advancing whisker based navigation through the implementation of bio-inspired whisking strategies," in *2016 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, Dec. 2016, pp. 767–773. DOI: 10.1109/ROBIO.2016.7866416. [Online]. Available: https://ieeexplore.ieee.org/document/7866416/ (visited on 06/30/2025).

[11] M. J. Pearson, B. Mitchinson, J. C. Sullivan, A. G. Pipe, and T. J. Prescott, "Biomimetic vibrissal sensing for robots," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 366, no. 1581, pp. 3085–3096, Nov. 12, 2011. DOI: 10.1098/rstb.2011.0164. [Online]. Available: https://royalsocietypublishing.org/doi/full/10.1098/rstb.2011.0164 (visited on 06/18/2025).

[12] C. Fox, M. Evans, M. Pearson, and T. Prescott, "Tactile SLAM with a biomimetic whiskered robot," in *2012 IEEE International Conference on Robotics and Automation*, May 2012, pp. 4925–4930. DOI: 10.1109/ICRA.2012.6224813. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/6224813 (visited on 07/02/2025).

[13] L. Andresen, E. Aucone, and S. Mintchev, "Whisker-based haptic perception system for branch detection in dense vegetation," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, Apr. 2022, pp. 911–918. DOI: 10.1109/RoboSoft54090.2022.9762143. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9762143 (visited on 06/13/2025).

[14] D. Jung and A. Zelinsky, "Whisker based mobile robot navigation," in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS '96*, vol. 2, Nov. 1996, 497–504 vol.2. DOI: 10.1109/IROS.1996.570842. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/570842 (visited on 05/20/2025).

[15] C. Ye, G. De Croon, and S. Hamaza, "A biomorphic whisker sensor for aerial tactile applications," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 5257–5263. DOI: 10.1109/ICRA57147.2024.10610850. [Online]. Available: https://ieeexplore.ieee.org/document/10610850 (visited on 04/30/2025).

[16] S. Hamaza, I. Georgilas, and T. Richardson, "2d contour following with an unmanned aerial manipulator: Towards tactile-based aerial navigation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov. 2019, pp. 3664–3669. DOI: 10.1109/IROS40897.2019.8968591. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8968591 (visited on 06/24/2025).

[17] C. Xiao, S. Xu, W. Wu, and J. Wachs, "Active multiobject exploration and recognition via tactile whiskers," *IEEE Transactions on Robotics*, vol. 38,

no. 6, pp. 3479–3497, Dec. 2022, ISSN: 1941-0468. DOI: 10.1109/TRO.2022.3182487. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9813357 (visited on 05/20/2025).

[18] M. A. Lin, H. Li, C. Xing, and M. R. Cutkosky, *Navigation and 3d surface reconstruction from passive whisker sensing*, Jun. 10, 2024. DOI: 10.48550/arXiv.2406.06038. arXiv: 2406.06038[cs]. [Online]. Available: http://arxiv.org/abs/2406.06038 (visited on 06/30/2025).

[19] T. Kent, "Whisker-inspired sensors for unstructured environments," thesis, Carnegie Mellon University, Mar. 7, 2025. DOI: 10.1184/R1/28513922.v1. [Online]. Available: https://kilthub.cmu.edu/articles/thesis/Whisker-Inspired_Sensors_for_Unstructured_Environments/28513922/1 (visited on 06/13/2025).

[20] Z. Yu, P. R. N. Childs, and T. Nanayakkara, *Towards the neuromorphic computing for offroad robot environment perception and navigation*, May 5, 2023. DOI: 10.48550/arXiv.2305.03860. arXiv: 2305.03860[cs]. [Online]. Available: http://arxiv.org/abs/2305.03860 (visited on 05/20/2025).

[21] S. Dey, D. Banerjee, A. M. George, A. Mukherjee, and A. Pal, "Efficient time series classification using spiking reservoir," in *2022 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2022, pp. 1–8. DOI: 10.1109/IJCNN55064.2022.9892728. [Online]. Available: https://ieeexplore.ieee.org/document/9892728 (visited on 11/19/2025).

[22] J. K. Eshraghian, M. Ward, E. O. Neftci, *et al.*, "Training spiking neural networks using lessons from deep learning," *Proceedings of the IEEE*, vol. 111, no. 9, pp. 1016–1054, Sep. 2023, ISSN: 1558-2256. DOI: 10.1109/JPROC.2023.3308088. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10242251 (visited on 04/24/2025).

[23] G. Brayshaw, B. Ward-Cherrier, and M. J. Pearson, "A neuromorphic system for the real-time classification of natural textures," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 1070–1076. DOI: 10.1109/ICRA57147.2024.10610401. [Online]. Available: https://ieeexplore.ieee.org/document/10610401 (visited on 06/05/2025).

[24] X. Xu, N. F. Lepora, and B. Ward-Cherrier, "A neuromorphic tactile system for reliable braille reading in noisy environments," *IEEE Robotics and Automation Letters*, vol. 10, no. 6, pp. 5225–5232, Jun. 2025, ISSN: 2377-3766. DOI: 10.1109/LRA.2025.3558707. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10955194 (visited on 06/05/2025).

[25] R. A. Russell and J. A. Wijaya, "Recognising and manipulating objects using data from a whisker sensor array," *Robotica*, vol. 23, no. 5, pp. 653–664, Sep. 2005, ISSN: 1469-8668, 0263-5747. DOI: 10.1017/S0263574704000748. [Online]. Available: https://www.cambridge.org/core/journals/robotica/article/recognising-and-manipulating-objects-using-data-from-a-whisker-sensor-array/121F1D0AB2031AEE4E845642F96982D7 (visited on 06/13/2025).

[26] W. Deer and P. E. I. Pounds, "Lightweight whiskers for contact, pre-contact, and fluid velocity sensing," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1978–1984, Apr. 2019, ISSN: 2377-3766. DOI: 10.1109/LRA.2019.2899215. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8641341 (visited on 05/29/2025).

[27] T. A. Kent, S. Kim, G. Kornilowicz, W. Yuan, M. J. Z. Hartmann, and S. Bergbreiter, "WhiskSight: A reconfigurable, vision-based, optical whisker sensing array for simultaneous contact, airflow, and inertia stimulus detection," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3357–3364, Apr. 2021, ISSN: 2377-3766. DOI: 10.1109/LRA.2021.3062816. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9366394 (visited on 06/02/2025).

[28] Z. Hu, Y. Cheng, J. Wachs, and Y. She, "Vibrissae-inspired vision-based magnetic-actuated whisker," *Nature Communications*, vol. 17, no. 1, p. 939, Dec. 21, 2025, ISSN: 2041-1723. DOI: 10.1038/s41467-025-67672-x. [Online]. Available: https://www.nature.com/articles/s41467-025-67672-x (visited on 02/09/2026).

[29] B. Mitchinson, J. C. Sullivan, M. J. Pearson, A. G. Pipe, and T. J. Prescott, "Perception of simple stimuli using sparse data from a tactile whisker array," in *Biomimetic and Biohybrid Systems*, Springer, Berlin, Heidelberg, 2013, pp. 179–190, ISBN: 978-3-642-39802-5. DOI: 10.1007/978-3-642-39802-5_16. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-642-39802-5_16 (visited on 06/17/2025).

[30] J. C. Sullivan, B. Mitchinson, M. J. Pearson, *et al.*, "Tactile discrimination using active whisker sensors," *IEEE Sensors Journal*, vol. 12, no. 2, pp. 350–362, Feb. 2012, ISSN: 1558-1748. DOI: 10.1109/JSEN.2011.2148114. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/5756632 (visited on 07/03/2025).

[31] D. Kim and R. Möller, "Biomimetic whiskers for shape recognition," *Robotics and Autonomous Systems*, vol. 55, no. 3, pp. 229–243, Mar. 31, 2007, ISSN: 0921-8890. DOI: 10.1016/j.robot.2006.08.001. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889006001400 (visited on 06/02/2025).

[32] P. Giguere and G. Dudek, "A simple tactile probe for surface identification by mobile robots," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 534–544, Jun. 2011, ISSN: 1941-0468. DOI: 10.1109/TRO.2011.2119910. [Online]. Available: https://ieeexplore.ieee.org/document/5752869 (visited on 04/10/2026).

[33] W. Maass, "Networks of spiking neurons: The third generation of neural network models," *Neural Networks*, vol. 10, no. 9, pp. 1659–1671, Dec. 1, 1997, ISSN: 0893-6080. DOI: 10.1016/S0893-6080(97)00011-7. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0893608097000117 (visited on 04/08/2026).

[34] C. Lv, Y. Wang, D. Han, X. Zheng, X. Huang, and D. Li. "Efficient and effective time-series forecasting with

spiking neural networks," arXiv.org. (Feb. 2, 2024), [Online]. Available: https://arxiv.org/abs/2402.01533v2 (visited on 01/08/2026).

[35] B. Ward-Cherrier, N. Pestell, and N. F. Lepora, "NeuroTac: A neuromorphic optical tactile sensor applied to texture recognition," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020, pp. 2654–2660. DOI: 10.1109/ICRA40945.2020.9197046. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9197046 (visited on 04/24/2025).

[36] H. Ismail Fawaz, G. Forestier, J. Weber, L. Idoumghar, and P.-A. Muller, "Deep learning for time series classification: A review," *Data Mining and Knowledge Discovery*, vol. 33, no. 4, pp. 917–963, Jul. 1, 2019, ISSN: 1573-756X. DOI: 10.1007/s10618-019-00619-1. [Online]. Available: https://doi.org/10.1007/s10618-019-00619-1 (visited on 04/10/2026).

[37] Q. Wen, L. Sun, F. Yang, *et al.*, "Time series data augmentation for deep learning: A survey," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, Aug. 2021, pp. 4653–4660. DOI: 10.24963/ijcai.2021/631. arXiv: 2002.12478[cs]. [Online]. Available: http://arxiv.org/abs/2002.12478 (visited on 02/09/2026).

[38] M. Yao, O. Richter, G. Zhao, *et al.*, "Spike-based dynamic computing with asynchronous sensing-computing neuromorphic chip," *Nature Communications*, vol. 15, no. 1, p. 4464, May 25, 2024, ISSN: 2041-1723. DOI: 10.1038/s41467-024-47811-6. [Online]. Available: https://www.nature.com/articles/s41467-024-47811-6 (visited on 06/12/2025).

3

Literature Study

Aerial Tactile Perception Using Whisker-Inspired Sensors and Spiking Neural Networks

A Literature Review

T.J. van der Weij

Technische Universiteit Delft

TU Delft

Aerial Tactile Perception Using Whisker-Inspired Sensors and Spiking Neural Networks

A Literature Review

by

T.J. van der Weij

in partial fulfilment of the requirements for the degree of
Master of Science
at the Delft University of Technology

Programme:	Robotics	
Department:	Cognitive Robotics	
Faculty:	Mechanical Engineering	
Student number:	4858999	
Project duration:	April 22, 2025 – July 9, 2025	
Supervisory team:	Asst. Prof. dr. ir. S. Hamaza	TU Delft
	Asst. Prof. dr. N. Tömen	TU Delft
	Asst. Prof. dr. ir. L. Ferranti	TU Delft

TU Delft

Abstract

Unmanned Aerial Vehicles (UAVs) increasingly operate in environments where traditional perception systems, such as cameras, LiDAR, and GNSS, struggle due to occlusion, poor lighting, or signal loss. Biological systems, by contrast, rely on tactile sensing through whiskers (vibrissae) to navigate and interact with their surroundings. Inspired by this, whisker-based tactile sensing has emerged as a promising research direction in robotics. These sensors offer lightweight, compliant, and energy-efficient solutions for close-proximity perception and can complement existing modalities in unstructured or vision-denied settings.

In parallel, Spiking Neural Networks (SNNs) have gained attention as a biologically plausible and power-efficient computing paradigm. They process information using discrete spikes, allowing event-driven and low-latency computation ideally suited for real-time robotic applications. Literature shows significant progress in tactile sensing and neuromorphic processing individually; however, their integration, particularly for aerial tactile systems, remains largely unexplored.

Building on this gap, the proposed research investigates how accurately a hybrid Multi-Layer Perceptron–Spiking Neural Network (MLP–SNN) can classify object shapes based on signals from a biomimetic whisker sensor array. This study serves as a first step toward combining tactile sensing and neuromorphic computing for energy-efficient perception in autonomous aerial systems.

Contents

Acronyms

ANN Artificial Neural Network. 22–24, 28, 29, 31, 33, 36, 37

AP Action Potential. 22–25

BNN Biological Neural Network. 22

BPTT Backpropagation Through Time. 28, 33, 34, 36

GNSS Global Navigation Satellite System. 1, 7, 14, 20, 32, 35, 37, 38

LIF Leaky Integrate-and-Fire. 25, 26, 31, 33, 34, 38

MAV Micro Air Vehicle. 18

MLP Multi-Layer Perceptron. 38

RNN Recurrent Neural Network. 13, 28

SLAM Simultaneous Localization And Mapping. 14, 15

SNN Spiking Neural Network. 1, 3, 21–39

STDP Spike-Timing Dependent Plasticity. 26, 29, 31

UAV Unmanned Aerial Vehicle. 1–3, 7, 12–14, 18, 20, 21, 31, 35–39

1 Introduction

1.1. Research Context and Motivation

Imagine perceiving the world primarily through touch. Many animals do exactly that by using their *vibrissae* (also known as whiskers) to explore and navigate complex environments through tactile interactions [1]. This sensory modality plays a critical role in perception, enabling the precise detection of nearby objects, obstacles, and textures [60].

In contrast, Unmanned Aerial Vehicles (UAVs) have traditionally been designed to avoid physical contact with their surroundings at all costs. However, nature demonstrates that contact and even collisions can be exploited to gain valuable information, providing animals with essential perceptual and navigational capabilities [35]. Shifting from a purely sense-and-avoid paradigm to a contact-based perception strategy could therefore expand the application potential of UAVs.

To fully exploit tactile sensing in aerial robotics, fast and efficient on-board processing is required. Neuromorphic computing provides a biologically inspired and energy-efficient framework for processing tactile information in real time [17]. In particular, a Spiking Neural Network (SNN) mimics the temporal dynamics of biological neurons, enabling low-latency decision-making in power-constrained platforms such as UAVs [17].

The combination of artificial tactile sensing and neuromorphic computing represents a promising research direction for robotic perception and navigation [64]. Physical touch can provide direct information about the presence of objects [63], as well as their properties such as shape, surface texture, stiffness, and weight [60]. Importantly, tactile perception complements visual and auditory sensing modalities, especially in environments where conventional, indirect sensors struggle to provide reliable data [63]. Moreover, tactile signals are inherently sparse, offering a low-bandwidth and low-latency source of information that can augment visual input [49].

In nature, tactile intelligence plays a central role in behavior. Seals, rats, and cats rely heavily on whiskers as a core sensory channel [1]. Humans also turn to touch in the absence of visual or auditory cues to interpret their surroundings [65]. In rodents, for example, facial whisker arrays are organized to enable the recognition of shapes and the discrimination of textures [35]. These animals actively control their whisker movements, a process known as active whisking, to maximize environmental contact and information gain [35]. Figure 1.1, adapted from [83], illustrates both biological and robotic examples of whisker use.

1.2. Problem Definition

UAVs are increasingly deployed in complex, unstructured, and low-visibility environments such as caves, forests, tunnels, and collapsed buildings. In such conditions, conventional sensing modalities including cameras, LiDAR, and Global Navigation Satellite System (GNSS) often underperform [45, 76, 62]. Visual systems are particularly unreliable in low-light, high dynamic range, or foggy environments, while GNSS signals are frequently unavailable indoors or underground [8]. Dense particulate matter, such as smoke or dust, can further degrade the performance of visual and laser-based systems, and even active acoustic or radar sensing may fail in cluttered or confined spaces [66].

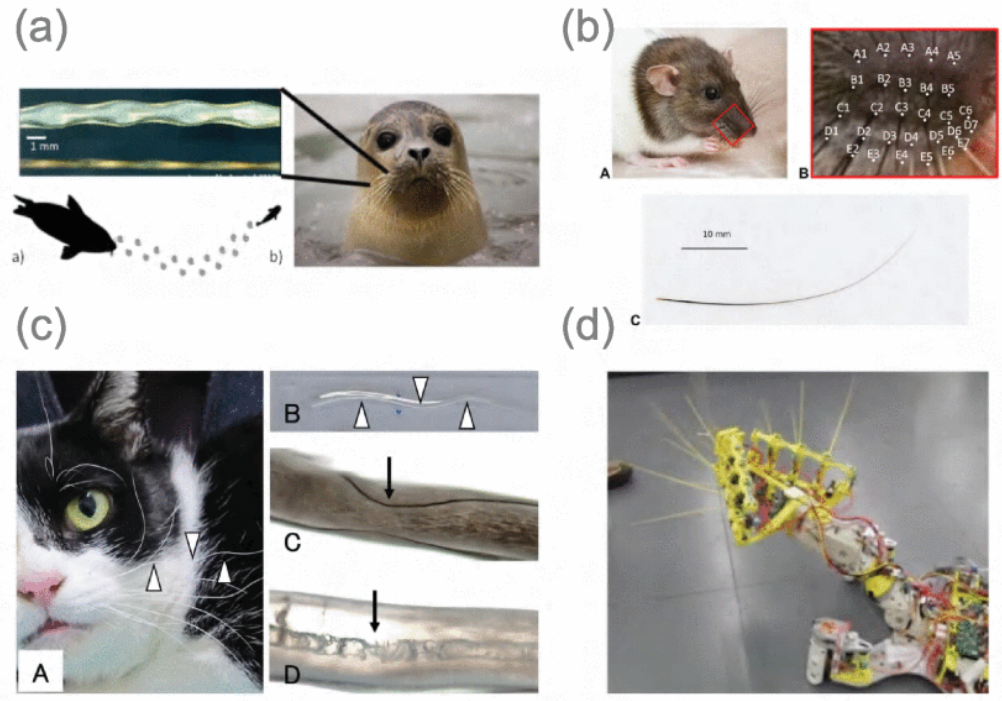

Figure 1.1: a) Seals use their whiskers to detect vortices underwater; b) Rats use their whiskers extensively for navigation and mapping; c) Cats rely on whiskers, among other functions, to catch prey; d) Artificial whiskers can provide tactile perception for robots. Adapted from [83].

These limitations highlight the need for alternative perception strategies in UAVs. Tactile sensing, unlike conventional methods, provides direct and robust physical feedback. Many animals rely on this modality to navigate dark or obstructed environments, using whiskers to detect obstacles and contours through contact [8]. This approach is both non-intrusive and energy-efficient, and it translates well to robotic systems [64].

Furthermore, UAVs aim to achieve maximum performance with minimal power consumption [81]. UAV platforms are inherently constrained by weight, power, and size, which limits the computational capacity of on-board hardware. At the same time, advanced tasks require energy-efficient, low-latency, and autonomous processing. Therefore there is a clear and pressing need for local processing solutions that combine high performance with strict payload and power consumption constraints [81].

Whisker-inspired sensors represent a promising tactile modality for UAVs. They are lightweight, flexible, and scalable, enabling reliable physical interaction with the environment even in conditions where vision fails [79]. When integrated into a multi-modal sensor suite, they can significantly enhance situational awareness and robustness [45, 76, 62].

The ability for UAVs to perceive and safely navigate in complex environments through touch has broad implications. Aerial tactile sensing is especially valuable in:

- Navigation and mapping of unstructured environments: autonomously exploring confined spaces (e.g. caves) where visibility is limited and physical contact is likely [40, 39].
- Search-and-rescue operations: locating survivors in collapsed buildings, where tactile navigation can play a critical role in life-saving missions [64].
- Covert and military applications: enabling stealthy operations with passive, low-power sensing strategies [55, 19].
- Environmental monitoring: documenting biodiversity in dense canopies or other environments inaccessible to vision-based systems [5, 73].

In each of these cases, aerial tactile perception, enabled by whisker-inspired sensors and neuromorphic computing, offers a real-time, efficient, and biologically plausible solution to environmental perception.

1.3. Research Scope

This literature review examines current scientific work on aerial tactile perception and neuromorphic computing, with a particular focus on whisker-inspired sensors. The guiding research question is:

> *What are the emerging research directions in aerial tactile perception using whisker-inspired sensors and Spiking Neural Networks?*

This question frames the exploration of recent advancements in tactile sensing, neuromorphic processing, and their integration for enhancing UAV capabilities in complex environments. To scope and structure the review, the following sub-questions (SQs) have been formulated:

- **SQ1**: What is tactile sensing, and what types of sensors are available, and what are its applications?
- **SQ2**: What are biological whiskers? And how have they inspired tactile sensing in robotics?
- **SQ3**: What are Spiking Neural Networks, and how do they differ from Artificial Neural Networks?
- **SQ4**: How can whisker-inspired sensors complement existing UAV sensors such as camera's and LiDAR?
- **SQ5**: How can tactile data from whisker-inspired sensors be processed using neuromorphic computing?
- **SQ6**: Why is tactile sensing promising for UAVs in low-visibility or unstructured settings?
- **SQ7**: Why is neuromorphic computing suitable for real-time tactile perception?
- **SQ8**: What are the potential societal, industrial, or environmental impacts of UAVs equipped with tactile sensing and neuromorphic processing?
- **SQ9**: Who could benefit from UAVs with tactile sensing?
- **SQ10**: Where are such UAV systems most needed?

In the introduction, SQ6–SQ10 are briefly considered to provide context. The main research question will be revisited in the conclusion, while SQ1–SQ5 form the core of the main outline of the report.

1.4. Organization

Section 2 presents the methodology applied in this literature review. Section 3 explores the domain of aerial tactile perception, focusing on tactile sensors and their applications, with particular attention to whiskers in both biological and robotic contexts. Chapter 4 introduces neuromorphic computing with SNNs, highlighting its distinctions from conventional neural networks. Chapter 5 summarizes the key findings of the review, addresses gaps in the literature and presents future reseach directions. Finally, Chapter 6 outlines the proposed master's thesis project.

2

Literature Review Method

For this research, a systematic and comprehensive literature search was conducted. A search plan was developed to ensure a comprehensive exploration of relevant topics and to identify high-quality sources for the study.

The process began with the creation of a general concept map (Figure 2.1), which provided an overview of the key themes and relationships within the research domain.

Based on this concept map, the main research question for the literature review was formulated as presented in the introduction. To address the research question, it was further analyzed into several terms and decomposed into sub-questions, which were discussed in Chapter 1 and are shown in Figure 2.2.

From these sub-questions, a structured search plan was designed. The most relevant concepts associated with the research topic are organized in a search table (Figure 2.3), from which specific search queries were derived (Figure 2.4). Alternative search terms were identified and combined using the Boolean operator OR, followed by integration of these groups with the operator AND. This strategy facilitated a targeted literature search across Google Scholar and Scopus, ensuring coverage of all critical aspects of the research topic.

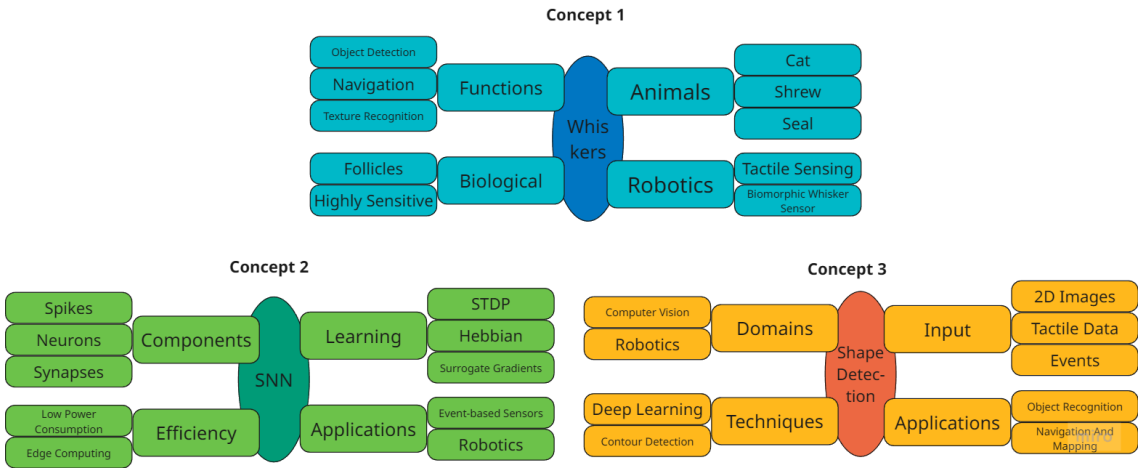

Figure 2.1: Concept maps of relevant concepts for the literature review.

Question	Explanation	Example
What?	**Definition of terms**	• **SQ1:** What is tactile sensing, and what types of sensors are available, and what are its applications? • **SQ2:** What are Spiking Neural Networks, and how do they differ from Artificial Neural Networks? • **SQ3:** What are biological whiskers? And how have they inspired tactile sensing in robotics?
How?	**Relation between terms**	• **SQ4:** How can whisker-inspired sensors complement existing UAV sensors such as camera's and LiDAR? • **SQ5:** How can tactile data from whisker-inspired sensor be processed using neuromorphic computing?
Why?	**Social relevance**: benefits of the research	• **SQ6:** Why is tactile sensing promising for UAVs in low-visibility or unstructured settings? • **SQ7**: Why is neuromorphic computing suitable for real-time tactile perception? • **SQ8:** What are the potential societal, industrial, or environmental impacts of UAVs equipped with tactile sensing and neuromorphic processing?
Who?	target group	• **SQ9**: Who could benefit from UAVs with tactile sensing?
Where?	geographical demarcation	• **SQ10**: Where are such UAV systems most needed?

Figure 2.2: Subquestions guiding this literature review, aligned with the main research question.

	Combine the concepts with AND		
Combine the synonyms with OR	whisker	spiking neural network	shape detection
	biomimetic whisker?	SNN	shape recogni*
	artificial vibrissa*	spiking neuron model	object recogni*
	vibrissa*	bio-inspired neural network	3D object detecti*
	artificial whisker?	event-driven neural network	geometric shape classif*
	tactile sensor?	neuromorphic computing	shape analysis
	whisker sensor?		3D contour mapping
	whisker array		

miro

Figure 2.3: Search Table for the literature review.

Search query 1	"whisker" OR "biomimetic whisker?" OR "artificial vibrissa*" OR "vibrissa*" OR "artificial whisker?" OR "tactile sensor?" OR "whisker sensor?" OR "whisker array"
Search query 2	"spiking neural network" OR "SNN" OR "spiking neuron model" OR "bio-inspired neural network" OR "event-driven neural network" OR "neuromorphic computing"
Search query 3	"shape detection" OR "shape recogni*" OR "object recogni*" OR "3D object detecti*" OR "geometric shape classif*" OR "shape analysis" OR "3D contour mapping"
Search query 4	[search query 1] AND [search query 2] AND [search query 3]

miro

Figure 2.4: Search Queries for the literature review.

3

Aerial Tactile Perception

Tactile perception refers to the interpretation of pressure data to gain information about the environment. In the context of aerial robotics, aerial tactile perception involves equipping platforms such as UAVs with the ability to sense and interpret contact or flow-based stimuli. This sense of touch can enhance a UAV's ability to navigate and explore environments that are challenging for traditional sensors like cameras or LiDAR.

Tactile sensing becomes particularly valuable in scenarios with poor visibility such as darkness, dust, fog, or cluttered indoor spaces and in GNSS-denied or unstructured environments [76, 8, 79]. While tactile sensing has been widely explored in ground-based robotics, its application in aerial platforms is an emerging area of research with distinct challenges and promising opportunities.

This chapter explores the concept of aerial tactile perception, beginning with a general overview of tactile sensing in robotics. It introduces different types of tactile sensors and discusses their relevance to UAV applications. Particular attention is devoted to whiskers from biological systems to robotic implementations, highlighting how their sensing principles can be adapted to aerial platforms. The chapter distinguishes between contact-based and airflow-based whiskers, each with their own applications and implementation challenges. Finally, it concludes with a summary of key insights and a transition to the next chapter on neuromorphic processing of tactile signals.

In addressing this topic, the chapter answers the following sub-questions: SQ1 - *What is tactile sensing, what types of sensors are available, and what are its applications?*, SQ3 -*What are biological whiskers, and how have they inspired tactile sensing in robotics?* and finally SQ4 - *How can whisker-inspired sensors complement existing UAV sensing systems such as cameras and LiDAR?*

3.1. Tactile Sensing in Robotics

In nature, tactile intelligence plays a central role in behavior. Seals, rats, and cats rely heavily on whiskers as a core sensory channel [1]. Humans also turn to touch in the absence of visual or auditory cues to interpret their surroundings [65]. In rodents, facial whisker arrays are organized to enable the recognition of shapes and the discrimination of textures [35]. These animals actively control their whisker movements, a process known as active whisking, to maximize environmental contact and information gain [35].

Humans rely on their sense of touch daily to perceive and interact with the physical environment. The human tactile system uses mechanoreceptors located beneath the skin to gather information about objects, such as their shape, texture, and compliance [29]. These mechanoreceptors efficiently encode tactile information, providing precise data on contact timing, location, and pressure [42]. Similarly, tactile sensing is crucial for robots to physically interact with their surroundings. It enables the perception of external stimuli, including surface roughness, textures, and shape, which are essential for tasks such as object recognition, manipulation, and navigation [42].

This section continues by presenting the various types of tactile sensors used in robotics, distinguishing between non-whisker-based and whisker-inspired approaches. This is followed by concrete examples of tactile sensors developed for robotic systems, laying the groundwork for understanding their relevance to aerial tactile perception.

3.1.1. Types of Tactile Sensors

To understand how robots acquire tactile information, it is essential to examine the different types of tactile sensors and the underlying transduction principles they employ. Tactile sensors are ”devices that acquire tactile information through physical interaction with the environment” [77]. They play a vital role in robotics by enabling systems to perceive and interact with their surroundings in a way that mimics biological touch. Tactile transduction techniques in such systems include piezoresistive, strain gauge, piezoelectric, capacitive, optical, and magnetic methods, each offering distinct advantages for different applications [80]. Yi, Zhang, and Peters provide a comprehensive review of the human sense of touch and its replication using biomimetic techniques [80]. Human mechanoreceptors, including Pacinian corpuscles, Meissner's corpuscles, Merkel discs, and Ruffini endings, enable recognition of shapes, textures, and hardness. These receptors exhibit adaptive behaviors classified as slowly or rapidly adapting [29]. Artificial tactile sensors emulate these biological properties, enabling robots to perform human-like tactile perception. These various sensing principles serve as the foundation for a wide range of tactile sensor implementations, which will be explored in the next subsection through specific examples used in robotic applications.

3.1.2. Overview of Existing Tactile Sensors

Building on the sensing principles outlined above, this subsection presents prominent examples of tactile sensors that have been developed for robotic applications, particularly those that incorporate biomimetic and neuromorphic design features. Notable examples include the BioTac sensor [18], which emulates the form and function of a human finger. It was developed to investigate tactile performance with the aim of approaching human-like perception. Building on this, Ward-Cherrier, Pestell, and Lepora developed the NeuroTac sensor, a novel neuromorphic optical tactile device that advances biomimetic tactile systems [74]. Neuromorphic tactile sensors mimic biological processes by employing spike-based representations inspired by the nervous system (see Chapter 4), thereby enabling more sophisticated tactile sensing inspired by human's sense of touch. In their work, hardware from the TacTip sensor [75] was combined with an event-based camera. Algorithms encoded contact information into spike trains for processing, enabling the NeuroTac sensor (Figure 3.1) to perform texture classification tasks and evaluate different coding methods' impacts on performance. A similar development is the NeuTouch fingertip sensor by Taunyazov et al., which contains 39 taxels (short for tactile pixels). This scalable sensor design supports larger numbers of taxels for robot end-effectors [72]. These examples illustrate how tactile sensors are evolving from conventional designs toward increasingly biomimetic and neuromorphic systems, paving the way for their integration into aerial platforms with advanced tactile perception capabilities.

3.2. Whiskers

Whiskers, also known as vibrissae, are specialized tactile hairs present in nearly all mammalian species [1]. The term originates from the Latin word *vibrio*, meaning “to vibrate” [1], which reflects their ability to oscillate and transmit mechanical stimuli to highly sensitive mechanoreceptors at their base. Vibrissae are integral to the sensory functioning and behavioral repertoire of animals, supporting activities such as navigation, prey detection, and environmental exploration in conditions where vision and hearing are limited. Whiskers are effective in confined and cluttered environments, especially where vision is limited [4, 40]. They provide valuable information about surrounding objects and surfaces as they make contact or brush past. Due to their low mechanical stiffness, whiskers can detect even very light contacts without disturbing the environment, a key advantage for state estimation [44].

This section first examines the biological foundations of whisker systems, including their morphology and the sensorimotor processes involved in tactile perception. It then explores how these biological principles have inspired the development of artificial whisker-based sensors. Distinctions are made between contact whiskers and airflow whiskers, each serving different perceptual roles. Finally, the section discusses practical applications of these whisker types and highlights recent robotic implementations that integrate them into real-world systems.

3.2.1. Biological Whiskers

To understand the principles underlying whisker-inspired tactile sensing, it is essential to first examine the biological function and adaptive behaviors of vibrissae in animals. In animals like cats and rats, vib-

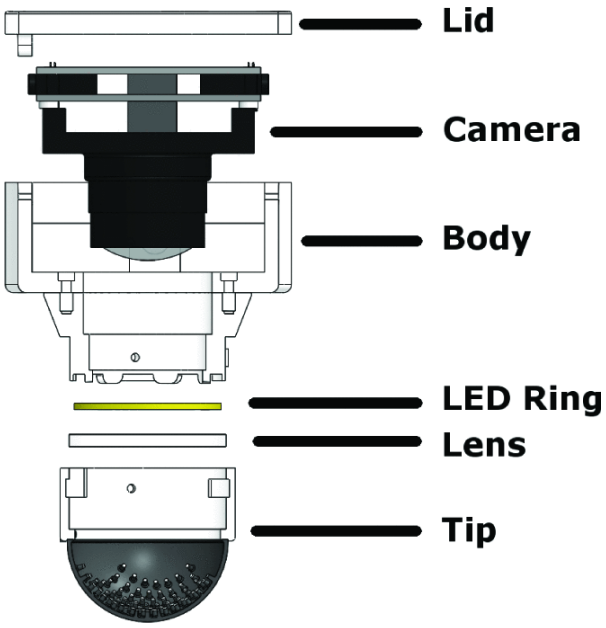

Figure 3.1: The recently developed NeuroTac sensor is a neuromorphic optical tactile sensor that enables the processing of event-based data for efficient and biologically accurate computing. Adapted from [74]

rissae supplement vision, enabling navigation in environments where visual cues are insufficient [76]. For instance, whiskers allow cats to assess whether openings are large enough for passage and detect minute changes in air currents deflected by nearby objects, facilitating movement in complete darkness without collisions [76]. Whisker-based touch forms a primary sensory modality for rodents, enabling them to detect floors, walls, and objects, especially in low-light conditions [14]. This sophisticated tactile system has inspired artificial whisker sensors in robotics, where similar sensing and processing capabilities are researched to enhance perception and autonomy. Figure 3.2 illustrates several mammalian species that rely on whiskers for environmental sensing [60]. Rats, seals, and shrews are particularly adept in using their whiskers. Whiskers naturally occur in various shapes. In Figure 3.3 an overview of different shapes of whiskers is presented, from [83].

While whiskers are often conceptualized as passive sensors, many animals actively control their vibrissae, a behavior named “active whisking.” As Prescott et al. note, active whisking offers several functional advantages:

1. Sampling across a broader spatial area.
2. Directing whiskers toward objects of interest.
3. Modulating contact force and duration for richer sensory input.

This dynamic sensing strategy enables animals to explore and recognize objects effectively in cluttered or dark environments. Its translation to robotics has inspired the development of actuated whisker arrays capable of similar exploratory and object recognition behaviors [49]. These biological insights form the foundation for designing artificial whisker systems.

3.2.2. Whisker Sensorimotor System

Beyond their mechanical properties, whiskers are part of a highly specialized sensorimotor system that allows animals to extract detailed spatial and textural information through neural processing. Research on whisker-based sensorimotor control and neuronal representations provides insights into how animals use vibrissae for navigation and object recognition. Diamond et al. investigated how ”where” (object location) and ”what” (object identity) are encoded in the rat’s neuronal system [14]. Their experiments included tasks such as navigating narrow alleys, discriminating cookie shapes, and recognizing textures. Rats were shown to use multiple whiskers simultaneously to infer object shapes, highlighting the importance of integrating contact information with self-motion cues. Non-intrusive sensing via

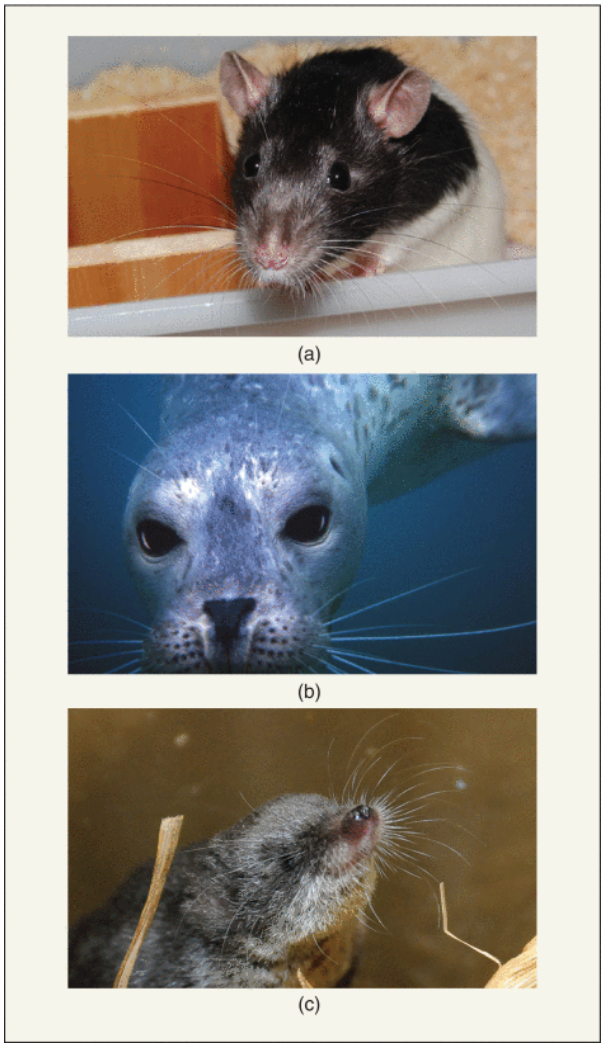

Figure 3.2: Examples of animals with whiskers. Three species with highly developed tactile sensory systems based on facial vibrissae: (a) common rat, (b) harbor seal, and (c) water shrew. Adapted from [60].

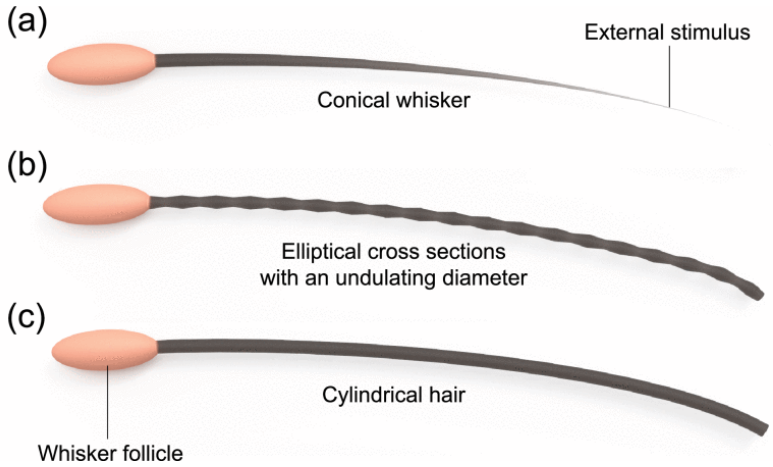

Figure 3.3: Overview of whisker shapes. (a) Rodents and large terrestrial mammals typically have conical whiskers. (b) Aquatic mammals, such as seals, possess whiskers with elliptical cross-sections and undulating diameters. (c) Cylindrical whiskers sensors are sometimes used in whiskered robots. Adapted from [83].

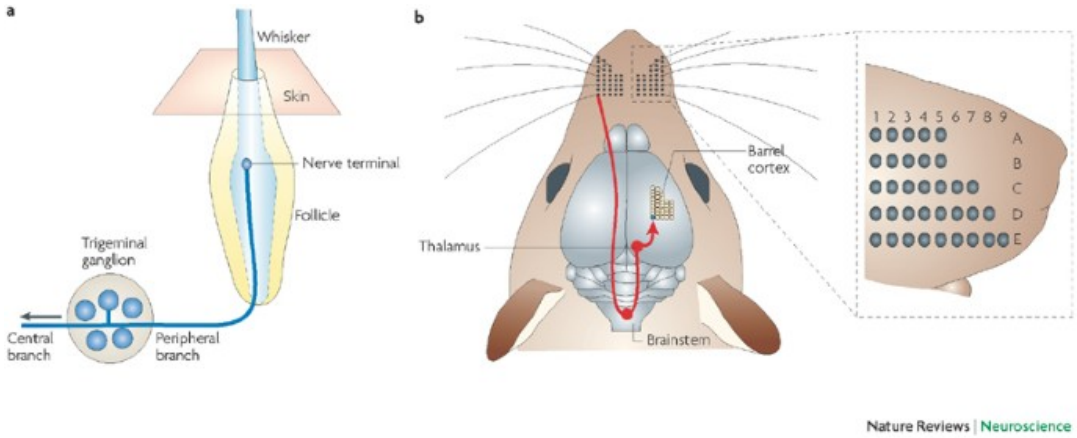

Figure 3.4: Schematic of the rat whisker sensory system illustrating the pathway from follicles to the barrel cortex. (a) Mechanoreceptors in each follicle respond to whisker deflections and rotations, encoding direction, velocity, and duration of stimuli. (b) Whiskers form a two-dimensional grid of five rows on each side of the snout. Adapted from [14].

whiskers is particularly advantageous for navigating and localizing objects in dynamic environments. In rats, the large whiskers (named macrovibrissae) are arranged in a precise grid of five rows (labeled A–E) and several numbered arcs on each side of the snout. Each whisker is rooted in a follicle richly innervated by approximately 200 sensory neurons of the trigeminal ganglion, which transduce mechanical deflections into action potentials [14]. These afferent signals travel through the trigeminal nerve to the brainstem's trigeminal nuclei, ascend to the thalamus, and ultimately project to the barrel cortex in the somatosensory region of the brain [14]. Remarkably, the layout of the whiskers on the snout is mirrored in the organization of the barrel cortex, forming a topographic "whisker-to-barrel" map (Figure 3.4) [14]. This highly ordered mapping facilitates the encoding of detailed spatial information about the environment. Research indicates that the direction, duration, and velocity of whisker deflections can be encoded in these neural signals [60]. Unique receptor responses in mammalian brains allow these animals to accurately estimate radial contact locations along their whiskers, inspiring similar approaches in robotics [61]. This intricate biological architecture demonstrates how mechanical deflections of whiskers are transformed into structured neural representations, insights that continue to influence the design of bio-inspired tactile processing in robotics.

3.2.3. Whisker-Inspired Sensors

Inspired by animal vibrissae, artificial whiskers have been developed as tactile sensors to enhance the perception capabilities of robots. Whiskers offer attractive features like high sensitivity, broad sensing range, mechanical compliance, robustness and simplicity in design [77, 67]. These engineered systems aim to replicate the unique properties of vibrissae that enable animals to gather rich spatial and textural information from their surroundings [77]. What makes whiskers particularly effective in the animal kingdom is the quality and quantity of sensory information they provide. They can detect subtle vibrations and surface textures without disturbing the environment, a feature especially valuable for state estimation in dynamic or delicate settings. Translating these capabilities into robotics allows for improved tactile exploration, object detection, and navigation in unstructured or visually impaired environments [83]. A comprehensive review of robotic platforms equipped with artificial whiskers is presented by Yu et al. in [83], highlighting various design approaches and applications across ground, aerial, and underwater robots.

Among whisker-inspired systems, the BIOTACT G2 sensor by Mitchinson et al. improves on the earlier G1 design [69]. It integrates 18 whisker modules, each with a composite tapered whisker mounted on a rotating shaft, motor, and controller for whisker protraction. Hall effect sensors measure protraction angles and two-dimensional whisker deflections [49]. Zhao, Jiang, and Li introduced a bio-inspired whisker array utilizing fiber Bragg grating sensors for measuring whisker deflection [85]. This lightweight, compact sensor offers high precision and robustness, supporting accurate 3D shape recognition in robotic applications. Optical tactile sensors also remain active in research. For example, the TacWhiskers system by Lepora, Pearson, and Cramphorn extends the TacTip family [75] by introduc-

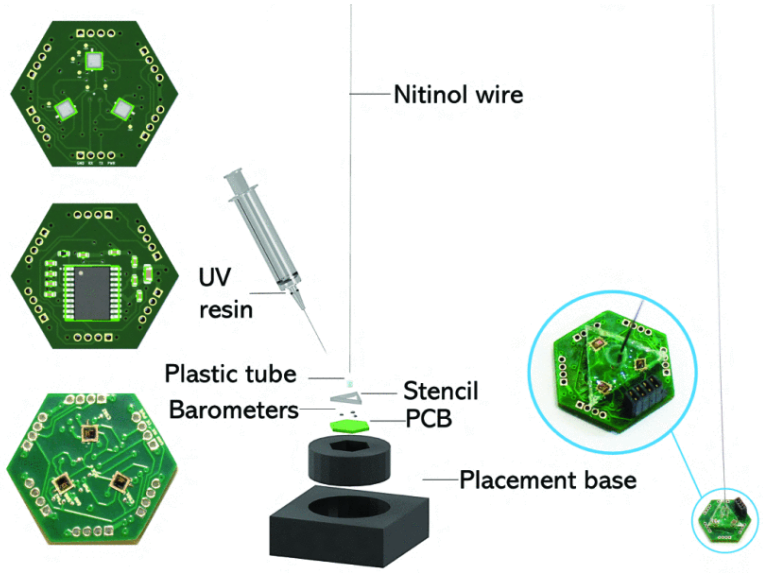

Figure 3.5: Biomorphic whisker design and fabrication process, featuring a nitinol wire, follicle structure, and integrated microcontroller PCB. This highly sensitive and scalable sensor can accurately localize contact points between the whisker and objects. Adapted from [79].

ing whisker arrays in static and dynamic configurations. The dynamic array allows controlled whisking motions, improving tactile perception accuracy. Designed for cost-effectiveness and customizability, TacWhiskers addresses issues of complexity and high costs in earlier tactile sensors [43]. Expanding on this approach, Kent et al. developed the Whisksight system which is a reconfigurable, vision-based whisker array sensor using a variation of the GelSight tactile sensor [84]. It measures whisker rotations and can separate contact, airflow, and inertial forces, supporting broader environmental sensing capabilities [33]. In [40], an innovative, open-source tactile sensor array was developed with telescoping whiskers adjustable in length and stiffness for task-specific optimization. This low-cost design supports navigation in confined spaces and compares favorably with LiDAR and cameras. The 16-whisker sensor demonstrated superior deflection sensitivity and range resolution, with decent spatial resolution and adaptability [40]. To address the limitations of traditional tactile systems, Xiao et al. developed a whisker tactile sensor with both high sensitivity and a wide sensing range, enabling more effective environmental perception [77].

In the recently published work by Ye, De Croon, and Hamaza [79], a lightweight, compact, and low-stiffness MEMS-based whisker sensor is introduced for integration on UAVs without compromising their flight performance. This design enables non-obtrusive interactions and offers a broad sensing range. The fabrication and structure of the whisker sensor are shown in Figure 3.5. To address key challenges in whisker sensor development, such as non-linearity, hysteresis, and the ambiguity of contact point localization (where multiple points can produce similar deflections), the authors proposed a novel approach that achieved high localization accuracy with low inference times. The sensor consists of a printed circuit board (PCB), barometric pressure sensors, and a nitinol wire serving as the whisker element. These lightweight whiskers are particularly advantageous for UAVs operating in dark, dusty, smoky, or turbulent environments, where conventional heavier sensors such as LiDAR may be ineffective [13].

These advancements in whisker-inspired tactile sensors demonstrate their growing potential to enhance robotic perception, particularly in environments where traditional sensing modalities fall short. In the next subsections, using robotic whiskers for either contact mapping or airflow sensing is discussed.

3.2.4. Contact Whiskers

Contact whiskers enable fine-grained tactile perception by identifying precise points of contact along the whisker shaft, a process known as radial contact estimation. By collecting multiple contact points, robots can reconstruct detailed 3D representations of nearby surfaces and objects [32]. Early work, such as Russell [63], demonstrated tactile whisker sensors for surface mapping using simple binary

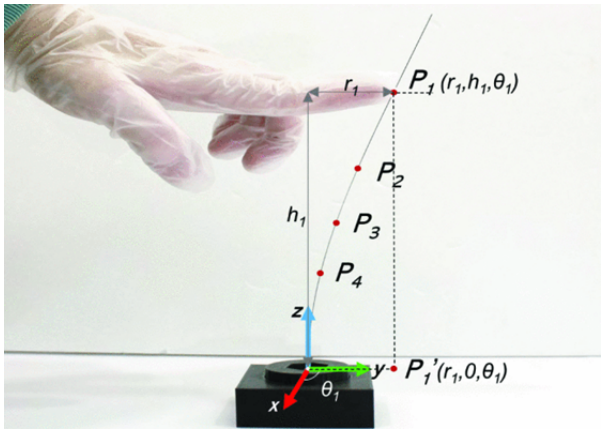

Figure 3.6: The 3D contact localization problem: Deflection at contact point P1 produces similar signal outputs as at P2, P3, and P4 due to the whisker's geometry. These points are shown in cylindrical coordinates. Adapted from [79].

contact/no-contact signals. However, accurately identifying the 3D contact location on a whisker has since emerged as a crucial requirement for efficient scene mapping [32]. At least nine different algorithms have been developed to estimate radial contact points on whiskers [32].

An essential aspect of tactile sensing with whiskers is thus not only detecting contact but also determining its precise location in space. As Wijaya and Russell notes, "If the contact point is at the whisker tip, then the information can be used directly to reconstruct the surface shape of the object. However, there is an ambiguity in the interpretation of the sensor data if the contact point is along the whisker length" [76]. This ambiguity complicates accurate surface reconstruction and requires additional processing strategies to resolve.

Moving from 2D to 3D localization further increases complexity. Robots must infer not only the radial distance but also the angular position and depth of contacts. This information is critical for reconstructing object shapes and supporting tasks such as manipulation, exploration, and navigation in unstructured environments [26]. However, multiple contact points along the whisker can produce nearly identical deflections, making it difficult to distinguish between them based solely on mechanical signals. To compute the 3D contact coordinates, algorithms typically rely on different types of sensing at the base(timing information, whisker strain, forces, or moments) and each approach comes with specific limitations [32].As shown in Figure 3.6, the deflection caused by contact at point P1 may generate similar signal outputs at subsequent contact points (P2, P3, P4), all defined in cylindrical coordinates [79].

One of the earliest algorithms for estimating contact points along a whisker was developed by Wijaya and Russell. The Tip Test Algorithm distinguished between tip and shaft contacts, using only tip contacts to reconstruct object surfaces with greater accuracy [76]. Building on this, Huet, Rudnicki, and Hartmann introduced the Elastica3D model, which maps six mechanical signals recorded at the whisker base to unique 3D contact point locations. Their findings showed that tapered whiskers outperformed straight ones in terms of localization precision, highlighting the importance of biomimetic design in whisker-based tactile systems [26].

Further simplifying the sensing process, Emnett, Graff, and Hartmann demonstrated that time derivatives of mechanical signals were not necessary for accurate localization. Instead, their approach used triplets of raw signals to directly estimate 3D coordinates [15]. Most recently, Ye, De Croon, and Hamaza addressed common challenges in whisker-based sensing, such as non-linearity and hysteresis, by developing a Recurrent Multi-output Network. This architecture integrates Recurrent Neural Networks (RNNs) to learn from streaming mechanical data and achieves high-accuracy 3D contact localization. Importantly, the sensor remains lightweight and scalable, making it well-suited for integration on UAVs operating in environments where traditional sensors like LiDAR are ineffective [79].

Addressing the complex challenge of radial contact estimation in whisker-inspired sensors can enable precise and reliable tactile perception in whisker-based robotic systems.

3.2.5. Airflow Whiskers

One promising application of whisker-inspired sensors is the detection of airflow [32]. In nature, vibrissae are not limited to contact-based sensing but also function as flow sensors. This capability is particularly developed in pinniped mammals, whose whiskers are both longer and denser with mechanoreceptors, featuring ten times more sensory receptors than those of terrestrial mammals [27]. These specialized whiskers allow marine animals to detect subtle water movements for tasks such as prey tracking and environmental navigation.

Translating this concept to aerial robotics, airflow sensing is particularly valuable for maintaining UAV stability in unsteady aerodynamic environments. As Hollenbeck et al. noted, airflow-sensitive tactile sensors can capture localized aerodynamic data, enabling rapid responses to disturbances such as gusts or turbulence [24].

Engineered airflow-sensing whiskers operate by deflecting under drag forces caused by either relative airflow (from the motion of the UAV itself) or global airflow (from wind in the environment) [32]. Unlike traditional sensors, whisker-inspired sensors are capable of detecting airflow in multiple directions, making them better suited for dynamic and unpredictable flight conditions [32].

With their multi-directional sensitivity and ability to distinguish between self-induced and environmental airflow, whisker-inspired flow sensors present an interesting research direction for enhancing situational awareness and flight control in UAVs.

3.2.6. Applications of Whisker-Inspired Sensors

The following section explores applications of whisker-inspired sensors in navigation, mapping, texture classification, object recognition and flow detection.

Navigation

In the natural world, many nocturnal or visually impaired animals rely on touch rather than vision to navigate their surroundings. By physically sensing their environment, these animals can determine the position and orientation of nearby objects and move safely. Tactile navigation uses physical contact to estimate both object locations and the animal's position relative to them [8]. Traditional UAV navigation primarily relies on GNSS and Simultaneous Localization And Mapping (SLAM) methods for localization and mapping [39]. However, in environments where GNSS signals are unavailable or visual sensors are obstructed, the sense of touch can be explored to enable autonomous navigation.

One of the earliest studies on whisker-based mobile robot navigation was conducted by Jung and Zelinsky on the Yamabico robot [31]. This system utilized proportional whisker sensors to represent the robot's inclination to walls and demonstrated high-speed wall-following capabilities. Two decades later, more advanced scientific research on tactile navigation was conducted. In 2019, Hamaza, Georgilas, and Richardson presented a UAV capable of contour following through tactile feedback using an energy-tank-based force control strategy. This approach supports applications such as crack marking during inspections and search-and-rescue missions [22]. In another study, Papachristos, Khattak, and Alexis developed a bio-inspired UAV system and experimented with its use in a corridor like environment. However, instead of using a energy tank system like in the paper by Hamaza, Georgilas, and Richardson, the authors use integrated antennae for haptic feedback in GNSS-denied and smoke-filled environments [53]. This design enhanced collision tolerance and resilience in confined spaces. Kossas et al. explored whisker sensors for wall-following and pathfinding in such environments. Their work demonstrated, both in simulation and real-world experiments, how tactile feedback can complement traditional techniques when performance deteriorates [39]. Another example is the work of Bredenbeck, Santina, and Hamaza, where a UAV was equipped with a compliant robotic finger inspired by human morphology and sensitivity [8]. The system allowed the UAV to follow contours and generate waypoints based on inferred contact forces, enabling robust navigation in visually and acoustically challenging environments. This is illustrated in Figure 3.7.

In a separate study [44], Lin et al. explored the use of a soft, curved whisker attached to a robotic arm to investigate tactile-based navigation and partial surface reconstruction. One of the key motivations for using whisker sensors in this work was their ability to overcome limitations inherent to traditional sensing modalities such as light, laser, and ultrasonic sensors. As Lin et al. point out, light and laser-based systems often fail when faced with challenging surfaces. Transparent objects tend to be nearly invisible to these sensors, while highly reflective surfaces can cause significant signal noise and measurement errors. Specifically, narrow infrared laser beams suffer from large errors near the edges of

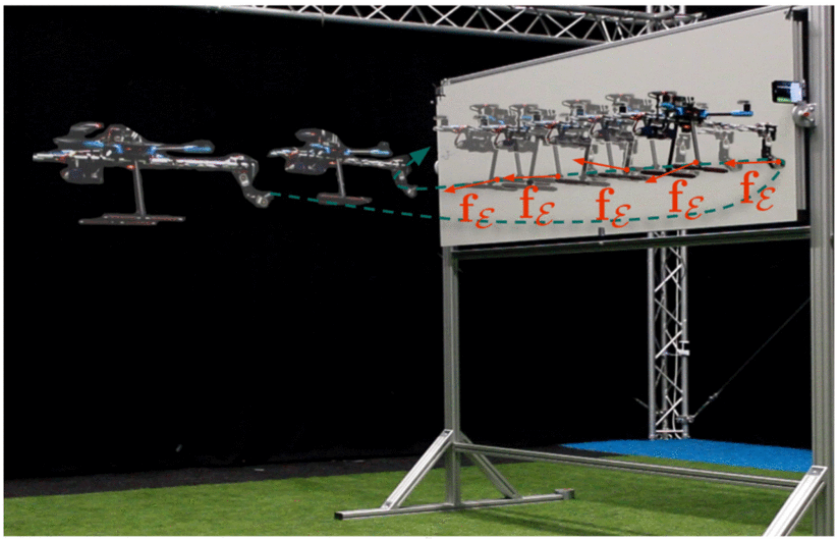

Figure 3.7: Aerial tactile navigation along walls with unknown shape and orientation: the quadrotor approaches the wall, makes physical contact using a compliant finger sensor, and computes new reference positions based solely on the deformation of the soft finger. Adapted from [8].

reflective objects [44]. Ultrasonic sensors, though more robust in some cases, exhibit high noise levels on irregular or non-flat surfaces due to their wide acoustic wave probes. By contrast, the whisker-based approach provides direct, physical interaction with objects, making it immune to such optical limitations. The tactile data collected from the whisker sensor array was used to generate occupancy maps, which are spatial representations of the environment that enable the robot to navigate and avoid obstacles effectively [44].

These studies highlight the potential of whisker-based tactile sensing to enhance UAV navigation in environments where traditional sensors fail, offering a physically grounded and resilient alternative for spatial awareness and obstacle avoidance.

Mapping

Tactile sensing also holds significant potential for applications in mapping, enabling robots to build maps of their surroundings through physical interaction. This capability is particularly valuable in scenario's where standard sensing methods face limitations. Whisker-based tactile sensing offers several advantages in this context: unlike vision systems, whiskers operate reliably in high dynamic range and are particularly effective at close distances [79]. Their compliant structure allows for gentle contact, reducing the risk of damaging either the sensor or the environment [32]. To enhance mapping speed and spatial resolution, arrays of whiskers are often employed, enabling robots to collect richer tactile data during exploration tasks [32].

Early studies, such as Fox et al., experimented with simple robots using four artificial whiskers for wall-following and exploration, laying groundwork for tactile SLAM systems [19]. In another example, Salman and Pearson applied biomimetic whiskers to unstructured environments. Building on the RAT-Slam algorithm [48], they demonstrated whisker-based mapping and localization capabilities [64]. In this bio-inspired approach, the aim is to develop a system that has the ability of "building a map of its environment and maintaining an accurate estimate of its location through whisker based touch" as stated by Salman and Pearson. Simultaneous estimation of both object pose and shape through tactile interaction is a complex problem [70]. To address this challenge, Suresh et al. proposed an online tactile SLAM framework for simultaneous object shape and pose estimation. A robotic arm equipped with a force/torque sensor applied controlled pushes to objects during multiple tasks, showing the feasibility of active tactile exploration in environments where visual sensing is inadequate [70].

These developments demonstrate how whisker-inspired systems can enable robust exploration and mapping, especially in environments where vision-based SLAM approaches are limited or unreliable.

Texture Classification
In general, tactile sensing for surface texture classification has been explored in several studies. For instance, Sullivan et al. demonstrated an early application using the first version of the BIOTACT sensor, which was mounted on a robotic end-effector. In their experiments, the robot performed surface texture recognition tasks by applying controlled vibrational patterns at specific frequencies. This approach enabled the system to distinguish between three different surface textures, showcasing the feasibility of biomimetic tactile sensing for robotic applications [69].

More broadly, such methods allow robots to detect and interpret fine-grained surface characteristics (such as roughness, smoothness, and patterns) that are critical for object manipulation and environmental interaction. Understanding surface properties can improve grip, handling precision, and the robot's adaptability to diverse materials. As Jamali and Sammut emphasize, integrating machine learning techniques with tactile sensors further enhances a robot's ability to classify and respond to surface textures in dynamic environments [29].

A related application is terrain classification for mobile robots, which also benefits from tactile sensing. Yu, Childs, and Nanayakkara developed a system using tapered whiskers as compliant springs, combined with reservoir computing, to reliably identify six distinct terrain types. The approach achieved high classification accuracy while maintaining low computational requirements, making it particularly suitable for resource-constrained platforms [81]. While demonstrated on ground robots, this method could be adapted to aerial platforms equipped with whisker-inspired sensors for in-situ surface identification.

These findings demonstrate that biomimetic tactile sensors, when paired with intelligent processing techniques, support both texture and terrain classification in complex and variable environments. They highlight the potential of whisker-based tactile systems for efficient, low-power, and accurate perception in aerial robotic platforms.

Object Recognition
Tactile-based object recognition has gained increasing interest as a robust alternative to vision in cluttered, low-light, or visually ambiguous environments [11]. Inspired by the sophisticated tactile perception of humans and animals, researchers have explored methods for robots to discriminate between object shapes and contours using whisker-like sensors.

Russell demonstrated in the 1990s that convex surfaces could be successfully measured using whisker sensors and proposed using arrays of whiskers for denser surface profile data [63]. Another early effort in this domain was conducted by Solomon and Hartmann, who developed a 4×1 array of robotic whiskers to reconstruct 3D object representations. By measuring the bending moments at the whisker bases, their system was able to infer the shape of contacted objects. They suggested this approach could significantly enhance robotic tactile perception and performance in unstructured environments [66]. Building on this, Solomon and Hartmann investigated two methods for inferring object shape: sweeping and tapping. Sweeping involved moving the whisker along or across an object to collect a series of contact points, while tapping used small rotational or translational movements to detect initial contact locations [67]. They developed a new algorithm for continuous object contour extraction using a single whisker sweep, which was tested on aluminium bars.

The influential work [77] of Xiao et al. focused on recognizing objects in unknown environments using a whisker-equipped UR16e robot. They proposed autonomous exploration strategies inspired by human tactile behaviors, such as contour tracing and volumetric sampling. During experiments (Figure 3.9), the robot's whisker sensor mounted on its end-effector explored cluttered scenes, successfully constructing object contours from contact points. Their deep learning classifier effectively categorized objects based on their contours (Figure 3.10), demonstrating strong performance across 11 real-world objects including a banana, hammer, and game controller [77]. These contours provided sufficient categorical information for accurate object classification. However, in highly cluttered environments, overlapping contours sometimes led to merged shapes, presenting challenges for recognition algorithms [77].

These advancements underscore the growing potential of whisker-based tactile systems, combined with learning models, to accurately recognize object shapes through physical interaction.

Flow Detection
Whisker-inspired airflow sensors on drones are a relatively recent innovation [32]. While their basic architecture consisting of a drag-sensitive body with sensing in the base shares similarities with bio-

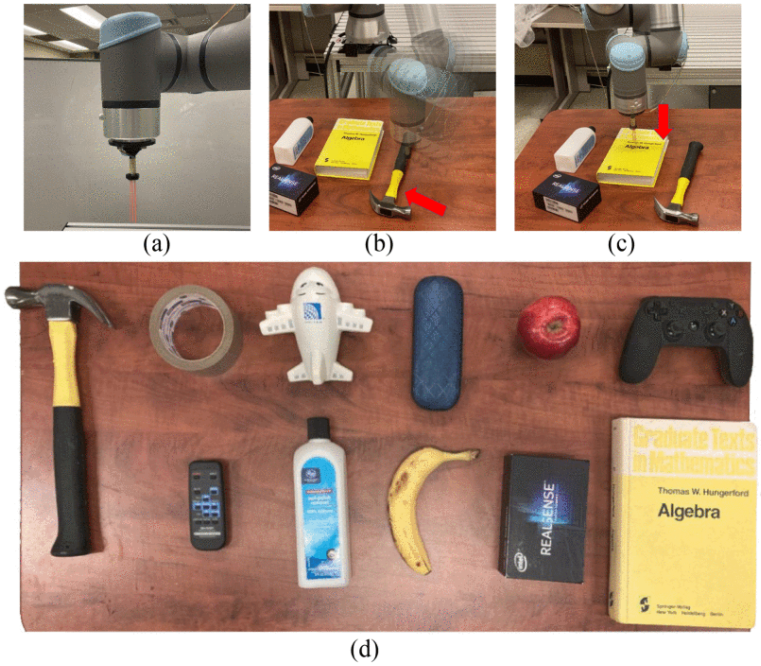

Figure 3.8: Whisker-based object detection process: contact point sampling and contour reconstruction. (a) Whisker sensor mounted on a UR16e robotic arm. (b) Object searching phase on a desktop environment. (c) Feature sampling by probing the objects from above to enhance shape discrimination. (d) The set of 11 real-world objects used in the experiments, including a banana, game controller, hammer, and apple. Adapted from [77].

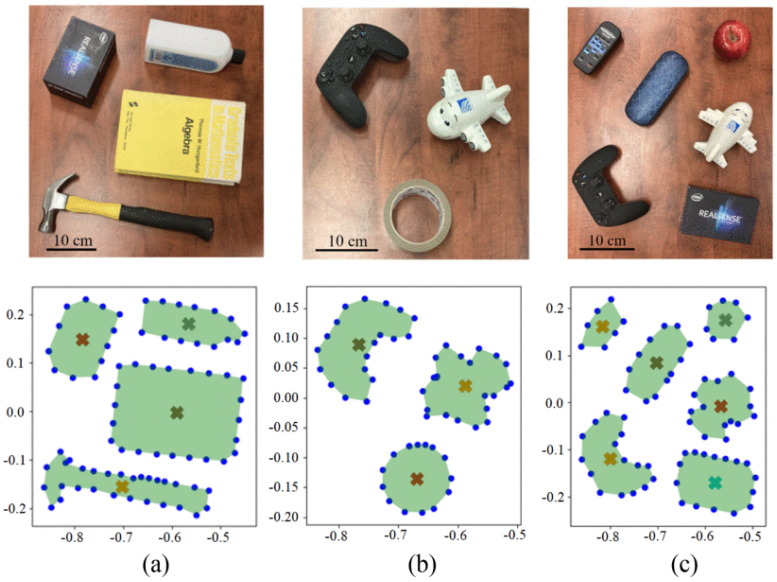

Figure 3.9: Tactile exploration experiments: scenes (top row) and reconstructed object contours (bottom row) with green regions representing recognized shapes. Adapted from [77].

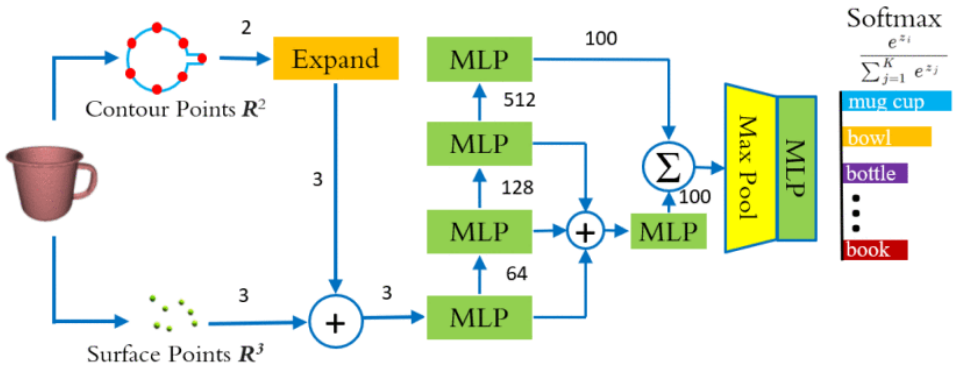

Figure 3.10: Deep learning object detection architecture that uses tactile contour and surface features to classify objects. Adapted from [77].

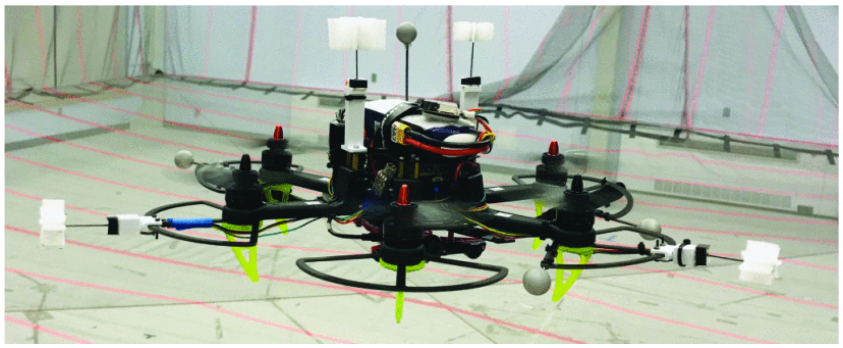

Figure 3.11: Airflow sensing on a multirotor equipped with four bio-inspired whisker-like sensors, enabling estimation of a 3D wind vector to distinguish aerodynamic drag from other forces. Adapted from [71].

logical vibrissae, engineered systems differ significantly in terms of symmetry, signal processing, and sensor integration [32].

In one example, Tagliabue et al. equipped a UAV with bio-inspired airflow sensors capable of distinguishing between airflow, drag, and interaction forces using deep learning models for relative airflow estimation [71]. Similarly, Deer and Pounds developed a lightweight whisker sensor designed for pre-contact force detection. This sensor enabled UAVs to perceive air pressure waves before direct contact, supporting improved obstacle avoidance and velocity control. Its low weight and robust construction make it suitable for aerial applications in dynamic environments [13].

A more complex design was introduced by Kim et al., who developed a whisker-fin sensor for multi-directional airflow detection. Inspired by animals that navigate using airflow cues, the design combines a whisker base with a fin structure to increase air resistance. Unlike conventional airflow sensors that are often limited to a single direction, this system enables detection across multiple directions, a critical feature for Micro Air Vehicle (MAV) stability in unpredictable airflow conditions. The sensor was validated both in wind tunnel experiments and real-world flights [37]. Related work on soft airflow sensors using artificial hair structures also shows promise for mobile robot platforms [38].

To address the challenge of distinguishing self-generated airflow from environmental wind, Kent and Bergbreiter developed a whisker sensor array that exploits a phenomenon they named ”flow shadowing” [34]. By arranging whiskers in a dense grid, airflow is unevenly distributed across the array, creating directionally dependent signal variations. According to Kent, this technique allows for the simultaneous detection of multiple flow sources by analyzing the signal response in relation to the flow heading. A secondary array can be used to resolve ambiguities between relative motion and ambient wind, enabling more precise control and interpretation of aerodynamic conditions.

These developments show that airflow-sensitive tactile sensors significantly enhance UAV responsiveness and flight control in complex aerodynamic environments, paving the way for safer, more adaptive, and robust flight strategies.

3.2.7. Whiskered Robot Projects

Robotic whisker systems have evolved significantly over the past two decades, drawing inspiration from the sophisticated tactile capabilities of animal vibrissae. Early implementations demonstrated the potential of whiskers for tasks such as navigation, object detection, and shape recognition, particularly in unstructured environments. This section provides an overview of complete robotic platforms that integrate whisker-based tactile systems, highlighting their design features, sensing strategies, and application domains.

One of the earliest platforms was Whiskerbot, introduced by Russell and Wijaya [62] and later refined in [57]. Whiskerbot featured an array of eight passive whiskers equipped with potentiometers to measure deflections. Since these whiskers were passive, the robot relied on its own motion to make contact with surrounding objects. Even at this early stage, researchers highlighted the benefits of tactile whiskers for navigation in poor visibility conditions, such as dust or low light.

The Koala platform, developed by Kim and Möller [35], advanced these ideas by implementing two arrays of biomimetic whiskers with Hall effect sensors and active whisking enabled by DC motors. Active whisking allowed the Koala robot to sweep its whiskers (steel beams with magnetic Hall-effect sensors) across objects, gathering tactile data to discriminate between simple geometric shapes (e.g., cubes, cylinders, and cones).

An improvement over Whiskerbot came with SCRATCHbot, developed by Pearson et al. [56]. This platform featured both micro- and macro-vibrissae and transitioned from potentiometers to magnets with Hall-effect sensors (like the Koala platform project), improving sensing performance and robustness.

Building on this work, the team introduced Shrewbot [55], which incorporated an array of 18 biomimetic whiskers that could be individually actuated. Hall effect sensors measured deflections, allowing the robot to explore surfaces and evaluate active whisking strategies. The design of the robot was inspired by the Etruscan shrew, the smallest terrestrial mammal, which exhibits remarkable tactile abilities. This animal can detect and capture prey, such as crickets, by using a single rapid whisk of its vibrissae to determine the prey's location and select an optimal bite point. Shape cues, like the cricket's legs, are critical in guiding its hunting behavior [14]. Experiments demonstrated how compliant, light-contact whisking improves tactile sensing and object recognition while minimizing disturbance to the environment. In more recent research, Andresen, Aucone, and Mintchev developed a whisker-based haptic perception system for navigating dense vegetation [5]. Unlike vision-based methods, this system exploited vegetation compliance to distinguish between rigid branches and flexible foliage, providing "pre-touch" awareness by detecting contact with whiskers before the robot itself reached an obstacle.

While whiskers have primarily been deployed on ground-based robots, recently scientists have researched their application to aerial and underwater platforms. For instance, Liu et al. designed whisker-inspired sensors to enhance underwater robots' perception, enabling obstacle avoidance and reactive behavior based on marine mammal whisker function [46]. Similarly, Guo et al. developed a piezoelectric whisker sensor for detecting underwater flow fields, inspired by seals' ability to sense subtle water disturbances caused by prey [21]. Although non-ground-based implementations have begun to emerge (particularly in the context of flow detection, as discussed in the previous subsection) complete aerial and underwater whisker-equipped platforms remain limited. Nonetheless, these advances demonstrate the growing potential of whisker-based systems in domains where conventional sensors often underperform, including aerial navigation and underwater exploration.

A non-exhaustive overview of robotic whisker projects is summarized in Table 3.1. Together, these projects demonstrate the growing versatility of whisker-inspired sensing systems, with ongoing advancements pushing their applications beyond ground-based robotics into aerial and aquatic environments where tactile perception offers distinct advantages.

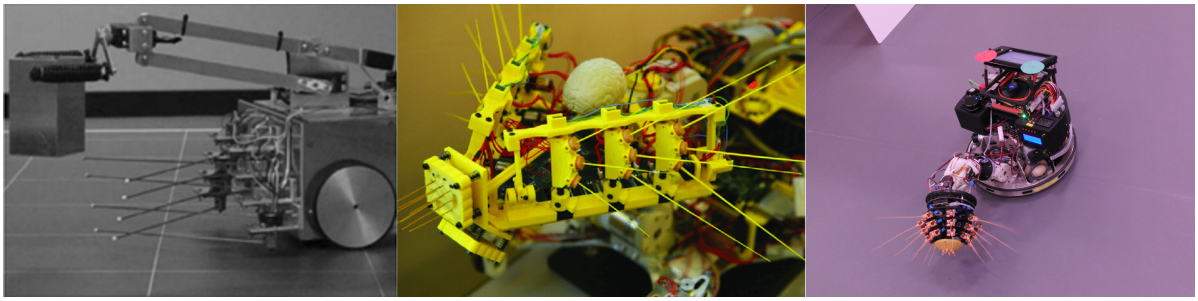

Figure 3.12: Whiskered robots, in chronological order from left to right: Whiskerbot 2005 (from [59], SCRATCHBot 2009 (from [50] and Shrewbot 2012 (from [10])

Whiskered Robot Project	Year	Description
aMouse	2004	A biomimetic whisker sensor system developed for a mobile platform, inspired by rodent vibrissae for active touch sensing.
Whiskerbot	2005	One of the first biomimetic robots using artificial vibrissae for texture discrimination, inspired by rat whiskers.
Koalabot	2007	A mobile robot integrating a small array of artificial whiskers for obstacle avoidance and surface following.
Scratchbot	2009	Designed for tactile exploration in complex environments, featuring a full array of artificial whiskers for navigation and surface characterization.
Whisking Sensobot	2011	Developed for haptic exploration in cluttered environments, using whiskers to sense object contours and textures.
Shrewbot	2012	Focused on using tactile sensing for object localization and navigation in total darkness, modeled after Etruscan shrews.

Table 3.1: Overview of whiskered robot projects and their key features.

3.3. Chapter Insights

This chapter explored the emerging domain of aerial tactile perception, focusing on the use of whisker-inspired sensors for UAV applications. It synthesized insights from biology, sensor design, and robotics to answer three guiding sub-questions.

- SQ1 — *What is tactile sensing, what types of sensors are available, and what are its applications?*

  Tactile sensing refers to the acquisition and interpretation of mechanical contact or pressure information. In robotics, tactile sensors have traditionally been used for manipulation and object recognition, and recent advancements have introduced biomimetic and neuromorphic designs to enable more sophisticated perception. Sensors vary widely in transduction methods including piezoresistive, capacitive, and optical techniques, and have been successfully applied to surface texture classification, shape recognition, and environmental exploration. For UAVs, tactile sensing offers an alternative in environments where vision or GNSS-based systems are unreliable or fail completely.

- SQ3 — *What are biological whiskers, and how have they inspired tactile sensing in robotics?*

  Biological whiskers (vibrissae) are highly sensitive tactile hairs used by mammals for navigation, object detection, and prey tracking, especially under poor visibility. Their characteristics of being lightweight, fast and non-intrusive has inspired the design of artificial whiskers for robots. These robotic whiskers mimic mechanical compliance, contact sensitivity, and array-based sensing to replicate the capabilities of biological systems. They are increasingly being researched on both ground and aerial robots for close-range sensing tasks.

- SQ4 — *How can whisker-inspired sensors complement existing UAV sensing systems such as cameras and LiDAR?*

  Whisker-inspired sensors offer critical advantages in environments where traditional sensors struggle. Unlike vision or LiDAR, which can be limited by lighting, dust, fog, or surface reflectivity, whiskers enable direct physical interaction and are capable of perceiving objects in low-visibility or unstructured settings. They enhance UAV perception in close-proximity tasks such as wall-following, terrain mapping, texture recognition, and airflow sensing, serving as complementary or fallback sensors for aerial navigation.

However, not only advanced sensor hardware but also intelligent signal processing techniques are needed for real-time autonomous control and flight of UAVs. The next chapter builds on this foundation by exploring the role of neuromorphic computing, particularly SNNs, in processing tactile signals from whisker-inspired sensors. It investigates how biological computing principles can be leveraged to enable efficient, event-based, and real-time perception for aerial tactile systems.

4

Spiking Neural Networks

SNNs have emerged as a promising approach for low-power, mobile, and hardware-constrained applications [47]. Unlike traditional Artificial Neural Networks (ANNs), which encode and process information as continuous values, SNNs transmit and compute information through discrete spikes over time. This spike-based communication is inspired by biological neural systems, where Action Potentials (APs) serve as the primary means of transmitting information between neurons [47]. By mimicking these temporal and sparse coding principles, SNNs represent a paradigm shift in machine learning toward energy-efficient and biologically plausible computation.

Building on the contents of Chapter 3, tactile perception, particularly from whisker-inspired sensors, generates high-dimensional, time-varying signals. SNNs provide a suitable computational framework for processing such inputs efficiently, making them highly relevant for neuromorphic aerial systems. Their spike-based architecture naturally aligns with the dynamic and event-driven nature of tactile data, offering the potential for low-latency, low-power, and scalable processing.

This chapter is organized as follows. It begins by introducing the principles of SNNs and comparing them to ANNs. Next, it explores how continuous tactile signals are encoded into spike trains and how SNNs can be trained using both biologically inspired and gradient-based methods. The chapter then presents relevant datasets, software tools, and signal processing techniques that support neuromorphic tactile perception. Finally, it highlights recent applications of SNNs in aerial tactile sensing, including object recognition and whisker-inspired systems, concluding with a discussion of challenges and future research directions.

In addressing this topic, the chapter answers the following sub-questions: SQ2 - *What are Spiking Neural Networks, and how do they differ from Artificial Neural Networks?* and SQ5 - *How can tactile data from whisker-inspired sensor be processed using neuromorphic computing?*

4.1. Foundations of Spiking Neural Networks

SNNs represent a new generation of neural network models, inspired by the temporal coding and sparse communication strategies found in biological neural systems. Unlike traditional ANNs, which continuously process input data at every timestep, SNNs operate in an event-driven manner, transmitting information only when significant input triggers a response [17]. This asynchronous and sparse computation enables SNNs to significantly reduce energy consumption while maintaining the capacity for complex information processing [17].

These characteristics make SNNs particularly suitable for real-time, low-power applications such as mobile robotics and embedded systems. In the context of aerial tactile perception, where fast reaction times, energy efficiency, and high-dimensional sensory input are essential, SNNs offer a compelling computational framework. Their biologically inspired encoding mechanisms, learning rules, and network architectures aim to combine the computational power of ANNs with the adaptability and energy efficiency of the human brain [47].

A key insight from biological systems is the efficiency of Biological Neural Networks (BNNs), in which the majority of neurons remain inactive unless stimulated by relevant sensory events [47]. This inherent sparsity enables energy efficiency in neural processing. As shown in Figure 4.1, sensory

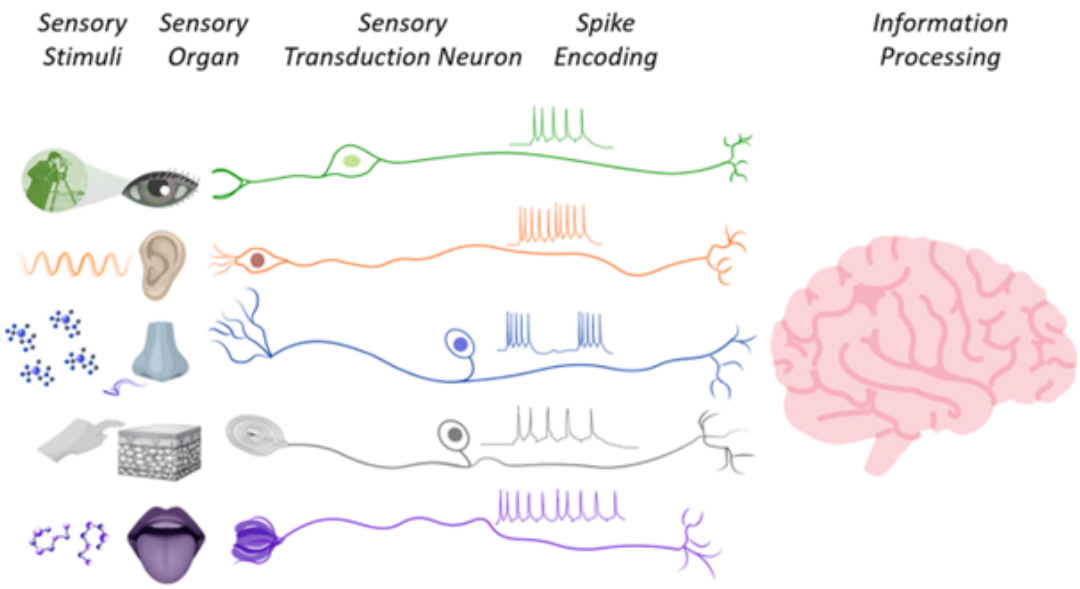

Figure 4.1: Schematic of a biological neural network (BNN). Sensory stimuli are perceived by specialized sensory organs, where neurons perform sensory transduction by converting analog inputs into spike trains. These spike trains are then relayed to the central nervous system for further processing. Adapted from [68].

stimuli are transduced into spike trains by specialized sensory neurons, which are then relayed through the central nervous system for further processing [68]. SNNs emulate this behavior by using spikes as the fundamental unit of communication, leading to computation that is both temporally precise and resource-conscious [17].

This section begins by presenting an overview of key terminology in neuromorphic computing. Next, SNNs are compared to conventional ANNs to highlight their distinctive properties and computational advantages. It continues by presenting an introduction to commonly used neuron models, ranging from biologically realistic to computationally efficient implementations. Finally, the complete neuromorphic pipeline, with neuromorphic sensors, algorithms and processors, is discussed.

4.1.1. Brief Overview

To understand how neuromorphic computing principles translate into artificial systems, it is first necessary to outline the biological foundations on which these systems are modeled, focusing on neurons, synapses, and action potentials. In biological systems, afferent neurons transform analog signals from sensory organs into spike trains. These spike-encoded signals are then transmitted to the central nervous system for processing [68]. Interestingly, biological transduction includes an element of stochasticity, which contributes to enhanced noise tolerance and energy efficiency [68]. In the following, a brief overview of neurons, synapses, and APs is presented. See Figure 4.2 from [17]for a visual overview.

- **Neurons:** Neurons are the fundamental atomic units of the brain, responsible for processing and transmitting information through electrical and chemical signals [47].
- **Synapses:** Synapses are the junctions between neurons, specifically the spaces between axon terminals of presynaptic neurons and the dendrites of postsynaptic neurons. They facilitate the transfer of information via neurotransmitters across these small gaps [47].
- **Action Potentials:** APs are discrete electrical signals, often referred to as ”spikes” through which biological neurons communicate [47]. These rapid events occur within milliseconds and can effectively be encoded as binary signals in computational models.

Together, these biological building blocks provide the basis for neuromorphic computing, enabling SNNs to replicate key properties of biological networks such as sparsity, event-driven signaling, and energy efficiency.

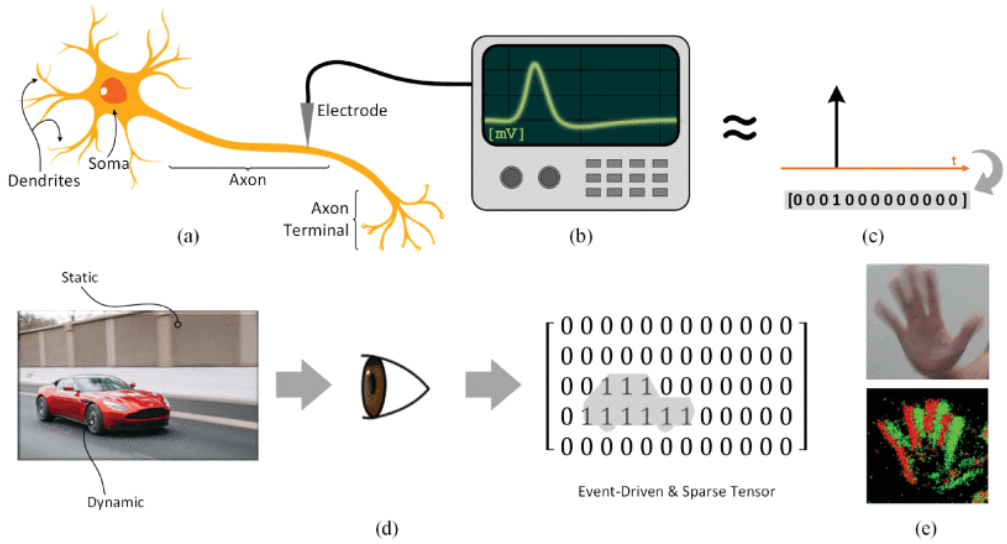

Figure 4.2: A visual overview of neurons and spikes. (a) Diagram of a biological neuron. (b) Measurement of an action potential propagating along the axon. (c) Approximation of a neuron's spike using a binary representation. (d) Event-driven processing. (e) Comparison between an active pixel sensor and an event-based camera. Adapted from [17].

4.1.2. Comparison with Artificial Neural Networks

Traditional ANNs encode and process information using continuous floating-point numbers. SNNs, on the other hand, communicate via discrete binary events. These spikes do not require precise floating-point representations, allowing SNNs to mimic biological neural processing more closely and potentially achieve greater energy efficiency [47].

The differences between ANNs and SNNs can be summarized in three key aspects according to Eshraghian et al.:

1. **Spikes:**
   Neurons in SNNs process and transmit information through discrete APs or spikes, whereas ANNs use continuous-valued activations [17].

2. **Sparsity:**
   Like biological neurons, which spend most of their time at rest, SNNs exhibit sparse activity. This means that at any given moment, most neurons are inactive, resulting in lower energy consumption and improved computational efficiency [17].

3. **Static Suppression:**
   SNNs are inherently more sensitive to dynamic changes in input, focusing on new or varying information. This mirrors biological sensory systems, which prioritize detecting changes in stimuli over static signals [17].

This comparison underscores how SNNs bridge the gap between biological realism and computational efficiency, positioning them as a promising paradigm for energy-efficient, event-driven processing in tasks such as tactile perception.

4.1.3. Neuron Models

Spiking neuron models define the computational core of SNNs, capturing the dynamic behavior of biological neurons through varying levels of abstraction and complexity. Artificial neuron models, as used in traditional ANNs, are relatively simple. Inputs are multiplied by corresponding weights, summed, and passed through an activation function (e.g., ReLU or sigmoid) to produce an output. In contrast, SNNs employ more biologically inspired neuron models. These models integrate weighted inputs over time to update the neuron's membrane potential. When the membrane potential exceeds a threshold, the neuron generates a spike and resets its potential.

To model this spiking behavior, researchers have proposed various neuron models of differing complexity, ranging from biophysically detailed frameworks like the Hodgkin-Huxley model to simplified abstractions such as the Leaky Integrate-and-Fire (LIF) model.

Leaky Integrate-and-Fire

The LIF model offers a simplified yet biologically inspired approach to modeling neuronal spiking, making it one of the most widely used neuron models in SNNs. Neurons accumulate incoming weighted inputs over time, updating their membrane potential $U(t)$. When the potential exceeds a threshold θ, the neuron fires and resets. As noted by Eshraghian et al., a spiking neuron closely resembles a low-pass RC circuit consisting of a resistor R and a capacitor C, leading to the LIF neuron model [17]. The dynamics of the passive membrane, modeled as an RC circuit, can be expressed as

$$\tau \frac{dU(t)}{dt} = -U(t) + I_{\text{in}}(t)R \tag{4.1}$$

where $\tau = RC$ is the membrane time constant, $U(t)$ is the membrane potential, and $I_{\text{in}}(t)$ is the input current. In discrete time, accounting for membrane potential decay and reset after spiking gives

$$U[t] = \underbrace{\beta U[t-1]}_{\text{decay}} + \underbrace{WX[t]}_{\text{input}} - \underbrace{S_{\text{out}}[t-1]\theta}_{\text{reset}} \tag{4.2}$$

where β represents the decay factor, $WX[t]$ is the weighted input at time step t, and $S_{\text{out}}[t-1]\theta$ subtracts the threshold θ from the membrane potential if the neuron fired at the previous time step. A spike is generated whenever the membrane potential exceeds θ, described as

$$S_{\text{out}}[t] = \begin{cases} 1, & \text{if } U[t] > \theta \\ 0, & \text{otherwise} \end{cases} \tag{4.3}$$

where $S_{\text{out}}[t] \in 0, 1$ indicates whether the neuron fires at time step t. Figure 4.3 illustrates a spiking neuron and the LIF model. In this model, both the input and output of the neuron are represented as binary spikes. The balance of biological plausibility and computational efficiency makes the LIF model a foundational building block for practical implementations of SNNs.

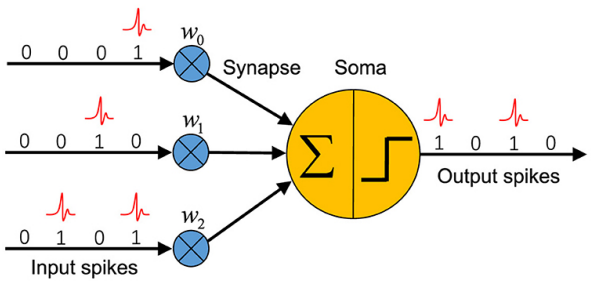

Figure 4.3: Spiking neuron model: schematic representation where both input and output are binary spike events. Adapted from [86].

Hodgkin-Huxley Model

A biophysically accurate and complex neuron model for the initiation and propagation of APs was introduced by Hodgkin and Huxley in 1952 [23]. The Hodgkin-Huxley model describes the neuron's membrane as an electrical circuit consisting of capacitors and resistors in parallel. The total current flow I in the neuron is given by:

$$I = C_m \frac{dV_m}{dt} + g_K(V_m - V_K) + g_{Na}(V_m - V_{Na}) + g_l(V_m - V_l)$$

where C_m is the membrane capacitance, V_m is the membrane potential, g_K, g_{Na}, g_l are conductances for potassium, sodium, and leak channels respectively, V_K, V_{Na}, V_l are their corresponding reversal potentials. This model treats the neuron's current flow as the sum of the capacitive current $C_m \frac{dV_m}{dt}$ and

the ionic currents through potassium (g_K), sodium (g_{Na}), and leak (g_l) channels. While highly detailed and accurate, the Hodgkin-Huxley model comes at a significant computational cost. For this reason, most SNNs use simplified models such as the LIF for practical applications.

Izhikevich Model
The Izhikevich model provides a trade-off between biological realism and computational efficiency [28]. It can reproduce many spiking behaviors of real neurons using a simple two-variable system:

$$\frac{dv}{dt} = 0.04v^2 + 5v + 140 - u + I$$

$$\frac{du}{dt} = a(bv - u)$$

with a reset condition:

$$\text{if } v \geq 30\,\text{mV}, \quad \begin{cases} v \leftarrow c \\ u \leftarrow u + d \end{cases}$$

where v is the membrane potential, u is a recovery variable, and a, b, c, d are parameters controlling the neuron's dynamics [28]. This model allows efficient simulation of networks with diverse neuron types and firing patterns. This flexibility enables the Izhikevich model to simulate large networks of heterogeneous neurons efficiently, making it well-suited for studying complex spiking dynamics in SNNs.

By offering a spectrum from biologically detailed to computationally efficient models, these neuron frameworks provide flexible tools for implementing SNNs to meet the specific application demands.

4.1.4. Neuromorphic Computing Pipeline

The neuromorphic computing pipeline integrates biologically inspired sensors, algorithms, and hardware to enable efficient, spike-based information processing in robotics. This pipeline comprises three key components: neuromorphic sensors, neuromorphic algorithms, and neuromorphic hardware [17].

Neuromorphic sensors generate event-driven outputs by encoding changes in sensory stimuli as asynchronous spikes, enabling efficient and low-latency perception. These sensors emulate biological sensory systems, enabling robots to interact with their environment in a more biologically plausible and responsive manner. Examples include event-based cameras and tactile sensors that directly generate spike trains suitable for processing by SNNs [74, 72].

Neuromorphic algorithms process spike-based data streams, utilizing temporal and spatial dynamics characteristic of biological networks. These algorithms leverage properties such as sparsity, asynchronous processing, and local learning rules (e.g., Spike-Timing Dependent Plasticity (STDP)) to achieve energy-efficient computation [47]. Applications include motor control, pattern generation, and navigation tasks like simultaneous localization and mapping (SLAM) [86].

Neuromorphic processors deviate from the traditional von Neumann architecture by employing massive parallelism, asynchronous event-driven operations, and highly distributed memory. This architecture reduces redundant computations and lowers energy consumption, making them particularly suitable for mobile and resource-constrained robotic platforms that require low-latency and energy-efficient processing [47, 86]. Architectures such as TrueNorth, Loihi, and SpiNNaker implement near-memory computing, leveraging sparse, event-driven processing for ultra-low power operation. However, most platforms are still limited to model inference and local learning and cannot yet train large-scale SNNs directly [86]. Table 4.1 summarizes currently popular neuromorphic chips.

Together, these three components form a complete pipeline that supports low-power, real-time, and adaptive computation, making neuromorphic systems especially promising for mobile and embodied AI applications.

4.2. Encoding and Training Spiking Neural Networks

SNNs process information through discrete spike events. Input data can either be fed directly as continuous values or transformed into spike trains using specialized encoding methods. Likewise, spike-based outputs must be decoded to yield meaningful predictions. A critical challenge in deploying SNNs is the training process, which is complicated by the non-differentiable nature of spikes.

Neuromorphic Chip	Developer	Characteristics
TrueNorth [3]	IBM	Digital neuromorphic processor with 1 million neurons, event-driven architecture, supports low-power inference.
Loihi [12]	Intel	Digital, supports local online learning (STDP, Hebbian rules), scalable, energy-efficient for SNNs.
SpiNNaker [52]	University of Manchester	Digital architecture designed for real-time large-scale neural simulations, massively parallel.

Table 4.1: Non-exhaustive overview of neuromorphic hardware platforms found in literature.

The section begins by introducing common input and output encoding techniques for transforming sensory data into spike trains and interpreting SNN outputs. It then presents four prominent training strategies for SNNs: shadow training, spike-based backpropagation, local learning rules, and surrogate gradient methods.

4.2.1. Input Encoding Methods

To effectively process continuous sensory data, SNNs may require input encoding methods that convert analog signals into spike trains, with different strategies offering trade-offs between efficiency, robustness, and temporal precision. Several encoding strategies have been proposed to achieve this conversion:

1. **Rate coding:** represents stimulus intensity as the firing rate of a neuron, where a stronger stimulus results in a higher number of spikes per unit time. This approach is robust to noise, making it reliable for many applications [17]. However, it suffers from higher power consumption and slower reaction times because it relies on accumulating spike counts over a period [17].
2. **Latency (or temporal) coding:** uses a time-to-first-spike model where only a single spike per neuron is required to encode the stimulus fully. Stronger stimuli result in earlier spikes, creating a negative correlation between stimulus intensity and spike timing [47]. This method offers reduced power consumption due to its sparse spiking behavior, which is particularly advantageous for low-energy systems [17]. However, it is more susceptible to noise and timing inaccuracies, which can affect the reliability of encoded information in dynamic environments [17].
3. **Delta modulation:** encodes temporal changes in input intensity by generating spikes whenever significant variations occur. This approach captures the derivative of the input signal, allowing for efficient compression of time-series data and making it well-suited for applications involving rapidly changing inputs [17]. While highly effective for such use cases, delta modulation's performance can degrade if the input signal is relatively static, as it may fail to produce sufficient spikes for accurate representation.

These encoding schemes (rate coding, latency coding and delta modulation) highlight the diverse approaches available for representing input stimuli as spikes.

4.2.2. Output Encoding Methods

The spike trains produced by SNNs must be decoded into meaningful outputs for interpretation. Common decoding strategies include:

1. **Rate coding:** The predicted class is determined by the output neuron with the highest firing rate (i.e., the most spikes within a time window) [17].
2. **Latency (or temporal) coding:** The prediction is based on which output neuron fires first, making it a fast and time-efficient strategy [17].
3. **Population coding:** A group of neurons collectively represents each class. The overall activity pattern across this population determines the predicted class [17].

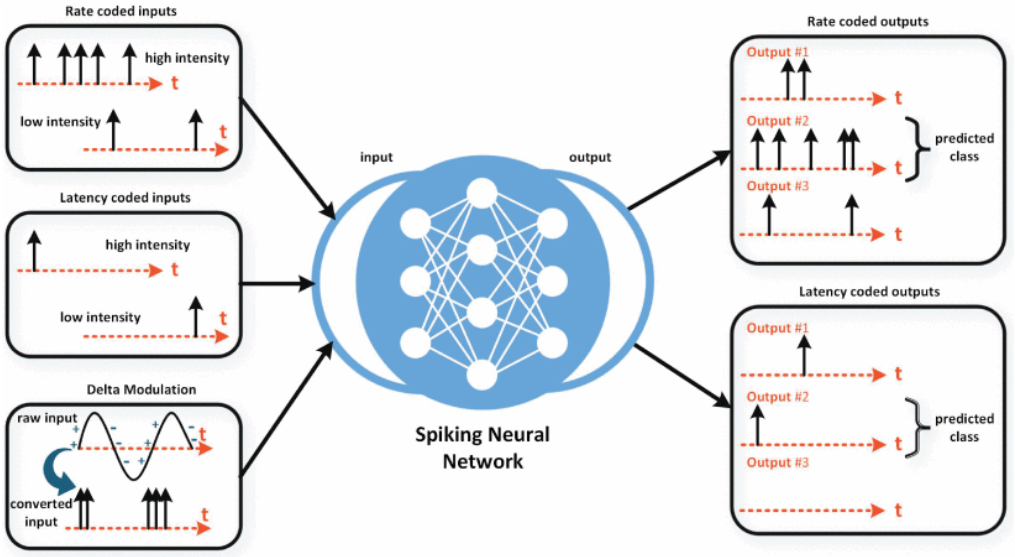

Figure 4.4: Visual representation of input and output encoding schemes in SNNs. Different methods such as rate coding, latency coding, and delta modulation illustrate how temporal changes in input intensity are transformed into spike trains for processing and how spike patterns are decoded for outputs. Adapted from [17].

These decoding methods (rate, latency and population coding) illustrate the flexibility of SNNs in transforming spike-based representations into actionable predictions or classifications.

In Figure 4.4, the various input and output encoding schemes used in SNNs are displayed visually, providing an overview of how continuous data can be represented as spike trains for processing [17].

4.2.3. Training Spiking Neural Networks

Training SNNs presents unique challenges due to the non-differentiability of spikes [17]. Unlike ANNs, where gradients can be computed directly through continuous activation functions, SNNs rely on discrete spike events, making traditional gradient-based optimization difficult [17]. Several approaches have been developed to address this. A common challenge in SNN training is the dead neuron problem, where neurons fail to fire due to poor weight initialization or suboptimal learning dynamics [17]. This results in underutilization of the network's capacity and can hinder performance.

Shadow Training

In shadow training, a non-spiking ANN is first trained using standard methods [17]. Once trained, the ANN is converted into an SNN by interpreting the neuron activations as firing rates or spike times. This approach takes advantage of the mature training algorithms and infrastructure of ANNs while enabling deployment in energy-efficient spiking architectures. However, this conversion process often leads to performance degradation due to differences in dynamics between ANNs and SNNs.

Backpropagation Using Spikes

An alternative is to train SNNs natively using error backpropagation, specifically Backpropagation Through Time (BPTT) [47]. This method propagates gradients through the temporal sequence of spike events, similar to how RNNs are trained. While computationally intensive, BPTT directly accounts for the temporal dynamics and sparsity of spike trains.

Local Learning Rules

Local learning rules update synaptic weights based on spatially and temporally local information at the synapse, rather than relying on a global error signal [47]. These biologically inspired methods include [17]:

- **Hebbian learning:** Synaptic weights strengthen when pre- and postsynaptic neurons are active simultaneously ("cells that fire together, wire together") [9].

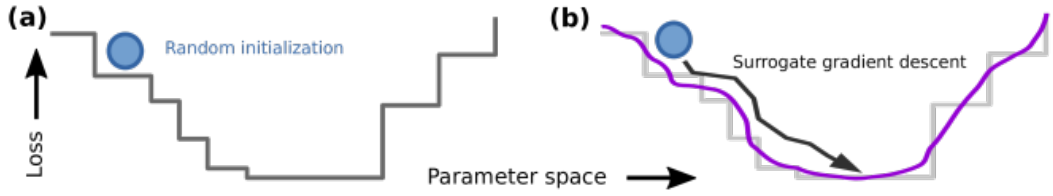

Figure 4.5: Surrogate gradients enable gradient-based training of SNNs despite the non-differentiability of spikes. Adapted from [47].

- **STDP:** A biologically plausible unsupervised learning rule where synaptic weight changes depend on the relative timing of pre- and postsynaptic spikes [47]. If a presynaptic spike precedes a postsynaptic spike, the synapse is strengthened; if the order is reversed, the synapse is weakened [17].

Surrogate Gradient Methods
To address the non-differentiability of spike functions, surrogate gradient methods replace the hard threshold of the spiking neuron with a smooth, differentiable approximation during training [51, 17]. This enables the use of standard gradient-based optimizers such as SGD or Adam. A schematic representation of surrogate gradient learning is shown in Figure 4.5.

Together, these training strategies, ranging from conversion-based to biologically inspired and gradient-based methods, highlight the diverse and evolving landscape of approaches for making SNNs trainable and effective in practical applications.

4.3. Tools and Resources for Neuromorphic Tactile Perception

Neuromorphic computing for tactile perception is an interesting research direction due to its potential for energy-efficient and low-latency sensory processing [78]. However, progress in this field is often hindered by the limited availability of suitable datasets and specialized tools [65]. This section introduces the publicly available neuromorphic tactile dataset ST-MNIST [65], and reviews software frameworks and processing pipelines that support the implementation and training of SNNs in tactile applications.

In this section, we progress from data acquisition to system implementation and, finally, to sensor interfacing and preprocessing. First, we review a software framework designed for building and training SNNs. Next, we present a neuromorphic tactile dataset that enables the benchmarking of SNN-based tactile systems. Lastly, we discuss recent advances in tactile signal processing, focusing on methods that translate analog sensor data into spike-based representations suitable for neuromorphic computation.

4.3.1. Software Framework

Software frameworks play a critical role in advancing the development and application of neural networks. Popular deep learning frameworks such as PyTorch provide powerful tools for constructing and training ANNs. [86] However, these frameworks are primarily developed for ANNs and do not natively support the unique requirements of SNNs. With the growing interest in large-scale brain-inspired neural networks, specialized frameworks have emerged to facilitate the modeling, simulation, and efficient computation of SNNs [86]. SNNTorch [17], a Python library, was developed to extend the capabilities of PyTorch and introduce networks of spiking neurons [16]. SNNTorch supports a range of biologically inspired learning algorithms, including variants of online backpropagation that are better aligned with the principles of neural computation observed in the brain. Building on this software framework, the next subsection considers an exaple of a neuromorphic tactile dataset to benchmark the performance of SNN models.

4.3.2. Tactile Dataset

Traditional tactile sensors often operate synchronously, using serial readout mechanisms to acquire data from multiple taxels. However, as the number of taxels or sensors increases, these synchronous systems face significant scalability challenges due to latency bottlenecks [65]. Event-based tactile sen-

sors, by contrast, process data asynchronously, enabling more efficient handling of large-scale sensor arrays and dynamic interactions. Despite these advantages, the lack of publicly available neuromorphic tactile datasets has historically limited progress in developing neuromorphic tactile systems. To address this gap, See et al. introduced a pioneering dataset: the Spiking Tactile MNIST (ST-MNIST) [65]. Analogous to the widely used MNIST dataset for vision tasks, ST-MNIST provides a benchmark for evaluating neuromorphic algorithms on tactile perception tasks. The ST-MNIST dataset consists of handwritten digits (0–9) collected from human subjects using a 100-taxel biomimetic event-based tactile sensor array. Each digit is recorded as a sequence of spatio-temporal events, capturing both the location and timing of contact across the sensor surface. This format closely mimics how biological systems perceive and process tactile information. Importantly, See et al. also reported baseline performances on this dataset. The introduction of datasets such as ST-MNIST provides a standardized benchmark for developing and evaluating neuromorphic tactile algorithms. While this datasets facilitates algorithmic development, the effectiveness of neuromorphic tactile systems also depends on how analog sensor signals are transformed into spike□based representations, which is discussed next.

4.3.3. Tactile Signal Processing

Biomimetic tactile perception with spikes has attracted increasing attentions in the past decade [80]. Advances in research can enable more biologically plausible and energy-efficient sensory systems. Inspired by the properties of biological mechanoreceptors, these approaches aim to replicate the temporal and spatial dynamics of natural tactile perception.

Over the past decade, research has explored how tactile sensors, typically producing analog output, can be interfaced with neuromorphic architectures through various analog-to-spike encoding techniques. Particularly promising are event-based tactile sensors, which bypass the need for explicit spiking neuron models by directly producing asynchronous spike trains in response to tactile stimuli [80]. These sensors offer a natural fit for spiking neural networks and represent a step toward fully integrated neuromorphic pipelines.

Table 4.2 provides an overview of studies in the last decade on biomimetic tactile perception using spikes. These works span a range of sensor types, encoding strategies, and decoding algorithms, with applications including gait event detection and texture discrimination. Despite their significance, many of these studies are now dated. More recent developments and application-focused studies will be discussed next.

Year	Author	Tactile Sensor	Analog-to-spike transformation	Spike train decoding	Application
2012	Spigler et al. [102]	A 2×2 piezoresistive tactile sensor array	The Izhikevich model	Spike frequency domain analysis	Grating discrimination
2013	Bologna et al. [91]	A 4×6 capacitive tactile sensor array	The integrate-and-fire model	Naïve Bayesian classifier	Braille letter recognition
2013	Lee et al. [104]	A fabric based binary tactile sensor array	The Izhikevich model	Tempotron	Curvature discrimination
2014	Lee et al. [135]	A low-cost foot pressure sensor	The Izhikevich model	Synaptic kernel inverse method (SKIM)	Gait event detection
2015	Rongala et al. [100]	A 2×2 piezoresistive tactile sensor array	The Izhikevich model	k-NN based on spike features	Texture discrimination
2015	Chou et al. [136]	Interactive neurorobot with 9×8 trackballs	The Izhikevich model	STDP	Learning touch preferences
2017	Yi and Zhang [112]	Biomimetic fingertip with PVDF films	The Izhikevich model	Spike train distance-based kNN	Roughness discrimination

Table 4.2: Tactile sensors and neuromorphic processing techniques. Adapted from [80].

Together, these datasets, software tools, and signal processing techniques form the foundation for implementing neuromorphic tactile systems. In the next section, we explore how these components are integrated in practice, particularly in the context of aerial robotics, where lightweight, energy-efficient, and responsive tactile sensing is essential.

4.4. Applications of Neuromorphic Tactile Perception

SNNs are increasingly adopted for tactile perception due to their energy efficiency, event-driven operation, and compatibility with asynchronous sensory data [72]. These features make them particularly attractive for mobile and aerial robots that operate under strict power and latency constraints. This section explores how SNNs have been implemented in neuromorphic tactile systems, object detection

tasks, and whisker-inspired perception architectures. These applications demonstrate the growing maturity and versatility of neuromorphic approaches in robotic sensing and perception, but also highlight future research opportunities.

In this section, we move from system-level implementations to specific use cases of SNNs in tactile perception. First, we present fully neuromorphic tactile sensing systems that integrate sensors, algorithms, and hardware for real-world tasks. Next, we highlight applications in object detection and recognition, showcasing how SNNs and hybrid architectures enable efficient and robust tactile perception. Finally, we examine whisker-inspired systems that explore biologically plausible models for processing tactile input.

4.4.1. Event-Based Tactile Sensing Systems

Neuromorphic tactile sensing systems combine event-based tactile sensors, spiking neural networks, and dedicated hardware to create fully integrated, low-power pipelines for real-world perception tasks.

In [78], a complete neuromorphic tactile sensing system is developed to tackle the real-world task of reading Braille letters. The system integrates a NeuroTac sensor [74] with a SCNN algorithm and deploys it on neuromorphic hardware, creating a fully neuromorphic pipeline. After training, the system achieves high classification accuracy on Braille letters. The setup is shown in Figure 4.6 (from [78]). Xu, Lepora, and Ward-Cherrier demonstrated that SNNs hold strong potential for increasing robustness to noise in dynamic and uncertain environments compared to conventional ANNs [78]. Ultimately, this work presents an energy-efficient tactile system with promising applications in low-power robotics.

In another study, Taunyazov et al. combine their NeuTouch tactile sensor with data from an event-based camera. By fusing visual and tactile information in a SNN, the system achieves improved classification performance. When deployed on Intel's Loihi chip, the SNN demonstrated comparable inference speed to an ANN while consuming approximately 1900 times less power [72]. This highlights the potential of end-to-end neuromorphic systems for power-constrained robotic platforms.

In [7], a fully neuromorphic system is presented for real-time texture classification using tactile data from the NeuroTac sensor. Similarly, Guo et al. propose DeepTactile, a novel network that leverages event-driven tactile sensing. By structuring tactile spikes from the NeuTouch sensor [72] as graphs and processing them with their model, the system achieves high classification accuracies while consuming up to ten times less energy compared to ANN baselines [20].

These studies demonstrate that fully neuromorphic pipelines (sensor + SNN + hardware) can operate with low energy consumption, real-time responsiveness, and noise robustness. These are all critical requirements for drones or UAVs that are power-constrained and operate in dynamic, uncertain environments (e.g., aerial inspection or search-and-rescue).

4.4.2. Event-Based Object Detection

Beyond tactile sensing, SNNs and hybrid architectures have been applied to event-driven object detection, making use of their asynchronous and energy-efficient processing for real-world tasks.

Regarding object detection, SNNs have particularly been applied in combination with event-based cameras. Kugele et al. propose a hybrid architecture that combines an SNNs backbone for efficient event-based feature extraction with an ANN head for classification and detection. This design uses the asynchronous, low-power processing of SNNs while maintaining high task performance [41]. In [25], the TacFormer network introduces attention mechanisms from Transformer models into SNNs for tactile object recognition. Evaluated on the EvTouch-Objects and EvTouch-Containers datasets, TacFormer demonstrates significant performance improvements over previous methods.

Another hybrid model is proposed in [2], which integrates SNNs and ANNs to combine spike-based processing efficiency with state-of-the-art detection performance. Using event camera data, the system achieves comparable accuracy to ANN baselines while reducing energy consumption and latency. Further contributions include Jin et al.'s R-SNN, which uses region proposals encoded as spike inputs for object detection [30]. Deployed on the Darwin Mouse neuromorphic system, R-SNN achieved improved energy efficiency, consuming only 6W during testing, though its performance still trails state-of-the-art ANNs [30].

In tactile sensing, Dabbous et al. address the limitations of fixed-rate sampling by developing a neuromorphic architecture with two layers of LIF neurons. Trained using STDP, their system achieves 100% accuracy in classifying eleven objects based on tactile shape data. Similarly, Kim et al. employ the Izhikevich neuron model to convert tactile signals into spike trains for unsupervised shape recognition.

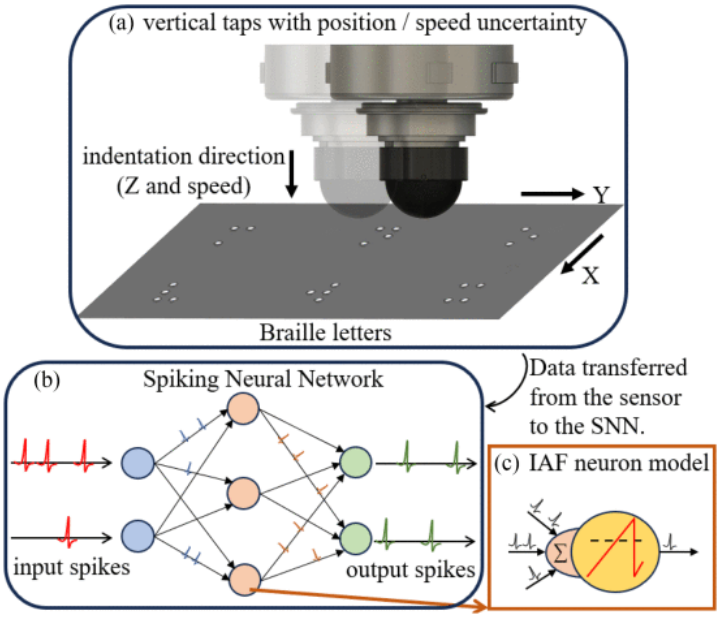

Figure 4.6: Neuromorphic Braille reading system using a NeuroTac sensor and SNN. (a) The NeuroTac sensor collects Braille information with inherent uncertainties in depth, center position, and speed. (b) Captured data is transmitted to the SNN for processing. (c) LIF neuron model employed in the system. Adapted from [78].

These applications highlight how SNNs can deliver competitive performance in real-world tasks like object detection while reducing energy and latency, emphasizing their potential for aerial tactile perception and other edge-computing scenarios. Event-based vision allows high-speed object detection with minimal power use, critical for drones that must detect and react to obstacles or objects in flight. In tactile contexts (e.g., drones physically interacting with surfaces), the use of SNNs for mapping, shape detection, or material recognition could enable intelligent touch-based interaction. These capabilities are essential for applications such as structure inspection, aerial manipulation, or autonomous navigation in unstructured environments.

4.4.3. Whisker-Inspired Sensors and Spiking Neural Networks

Moving on from event-based tactile sensing and object detection, we focus on whisker-inspired sensors. Building on the biological inspiration of vibrissae, several works investigate how spiking neural models can replicate the processing of whisker-based sensory input, providing insights into future neuromorphic tactile systems.

Although no studies have yet applied SNNs directly to whisker-based tactile systems, several works investigate how biological systems process whisker input. In [6] the authors develop a biologically inspired neurorobotic whisker system to replicate the functional mechanisms of animal somatosensory processing. Their SNN model processes information encoded during whisking, offering insights into neural dynamics. A related effort is presented by Petschenig et al. in [58], who focus on evoked sensory responses in the rat barrel cortex. Their work contributes to understanding how whisker sensory data can be processed in neuromorphic architectures. These approaches allow for better understanding of the neuroscientific background for merging bio-inspired tactile sensing and neuromorphic computing in tactile platforms.

Whisker-inspired sensors may be valuable in close-proximity sensing and collision detection in areas where vision and LiDAR are less effective due to occlusions or poor lighting (e.g., in tunnels, under bridges, or near infrastructure) [8]. The use of SNNs for whisker signal processing, even in simulation or neuroscience contexts, offers the potential for biomimetic, low-latency autonomous tactile navigation in aerial robots operating in tight or GNSS-denied spaces.

Collectively, these applications illustrate the versatility of spiking neural networks in tactile perception, from low-power sensing systems to object recognition, setting the stage for their integration into aerial platforms that demand efficient, real-time interaction with the environment.

4.5. Challenges in Spiking Neural Networks

Despite the promising advancements of SNNs in recent years, several key challenges remain that hinder their broader adoption and performance parity with ANNs.

- Firstly, although many studies have demonstrated that deep SNNs can achieve accuracy comparable to their ANN counterparts on certain tasks, they still lag behind conventional state-of-the-art ANNs, particularly on large-scale datasets such as ImageNet [86]. Bridging this performance gap requires further research efforts in network architecture design, training algorithms, and optimization techniques.
- Secondly, while SNNs are often praised for their energy efficiency, this metric does not account for the energy cost of other processes such as data movement and memory access [86]. To accurately assess energy savings, it is crucial to deploy well-performing SNNs on neuromorphic hardware or high-fidelity simulators that fully exploit the event-driven nature of these systems and provide realistic energy consumption measurements [86].
- Thirdly, applications requiring high processing speed and low power consumption, such as robotics, could greatly benefit from integrating neuromorphic vision and audio sensors with neuromorphic processing chips. However, current approaches often convert event-based sensor data into frame-based representations before processing, which undermines the advantages of event-driven systems. Greater efforts are needed to develop methods that directly process raw neuromorphic sensing events using SNNs, thus maintaining their inherent temporal resolution and energy efficiency [86].

Addressing these limitations is essential to fully realize the potential of SNNs in real-world applications, particularly in domains like robotics, where low-latency, energy-efficient, and event-driven processing can offer a substantial advantage over conventional approaches.

4.6. Chapter Insights

This chapter explored the foundations, implementation, and applications of SNNs, focusing on their relevance for tactile perception in aerial systems. By doing so, it aimed to answer two guiding subquestions.

- SQ2 — *What are Spiking Neural Networks, and how do they differ from Artificial Neural Networks?*

  SNNs differ fundamentally from ANNs by using discrete spikes instead of continuous activations to encode and process information. This event-driven approach mirrors the asynchronous signaling of biological neural systems, enabling energy-efficient and low-latency computation. Furthermore, SNNs exhibit sparse activity, meaning most neurons are inactive at any moment, which is unlike ANNs, where neurons are always 'on'. Finally, SNNs focus on new and dynamic changes, mirroring biological sensory systems. Compared to ANNs, SNNs can achieve similar task performance on certain benchmarks while offering advantages in temporal precision and scalability for embedded and mobile platforms.

- SQ5 - *How can tactile data from whisker-inspired sensor be processed using neuromorphic computing?*

  This chapter showed that tactile data, especially from high-dimensional, time-varying whisker-inspired sensors, can be efficiently encoded as spike trains and processed by SNNs using neuromorphic algorithms and processors. A spiking neura model captures the dynamics of biolpogical neurnosn and replaces the artiical neurons models, of which the most popular is the LIF model for LIF for practical implementations. Input encoding methods consist of rate coding, latency coding and delta modulation, wehreas output encoding skips deltat modulation and has rate, latenct and populstion coiding. Seeral training stragetfies exist, of wihc shadowrding traing, BPTT, local learning rules and suurgoate gradient methods are mentioned By employing encoding and training strategies for SNNs, such systems achieve low-power operation, robustness to noise, and real-time responsiveness, which are critical for aerial tactile perception.

  This chapter demonstrated that tactile data, particularly from high-dimensional, time-varying, whisker-inspired sensors, can be efficiently encoded as spike trains and processed by SNNs

using neuromorphic algorithms and hardware. Unlike traditional artificial neuron models, SNNs employ biologically inspired spiking neuron models, with the LIF model being the most commonly used for practical applications. Input encoding strategies include rate coding, latency coding, and delta modulation, while output encoding typically involves rate coding, latency coding, or population coding. Several training approaches have been developed to address the challenges posed by the non-differentiable nature of spikes, including shadow training, BPTT, local learning rules, and surrogate gradient methods. By combining appropriate encoding and training techniques, SNN-based systems can achieve low-power operation, robustness to noise, and real-time responsiveness, which are all key advantages for aerial tactile perception.

5

Conclusion

UAVs are increasingly deployed in complex, unstructured, and low-visibility environments such as caves, forests, tunnels, and collapsed buildings. In such settings, conventional sensing modalities (including cameras, LiDAR, and GNSS) often underperform due to occlusion, poor lighting, or signal loss [45, 76, 62]. At the same time, UAVs are inherently constrained by payload capacity, power availability, and size, limiting the computational capabilities of onboard systems. As UAV missions become more advanced and autonomous, there is a growing demand for perception systems that are not only low-power and low-latency, but also capable of functioning reliably under these adverse conditions [81]. These limitations underscore the need for alternative perception strategies that combine robust performance with stringent energy and hardware constraints.

This conclusion discusses the key findings from Chapter 3 and Chapter 4 in light of this research problem. First, we summarize the main insights of the literature review. Then, we critically reflect on the findings and identify gaps in the current literature. Finally, we offer a perspective on future opportunities by returning to the main research question: *What are the emerging research directions in aerial tactile perception using whisker-inspired sensors and Spiking Neural Networks?*

5.1. Summary of Findings

In this section, we summarize the key findings from the two main chapters: one on aerial tactile perception and the other on SNNs.

Aerial Tactile Perception

This chapter examined the emerging field of aerial tactile perception, highlighting how tactile sensing (particularly through whisker-inspired sensors) can enhance UAV capabilities in environments where traditional modalities like vision, LiDAR, or GNSS often underperform. Tactile sensing refers to the interpretation of mechanical interactions such as contact or pressure, with applications ranging from robotic manipulation and texture classification to environmental exploration (SQ1). A range of sensor types exists, including piezoresistive, capacitive, optical, and neuromorphic designs. Drawing inspiration from biology, whisker-inspired systems replicate the functionality of mammalian vibrissae, lightweight and sensitive hairs used for close-range perception in poor visibility. These artificial whiskers mimic mechanical compliance and contact-based sensing and have been implemented on both ground and aerial platforms (SQ3). The chapter reviewed key applications such as navigation, mapping, object recognition, and airflow sensing, showing that whisker-inspired systems serve as robust, low-power complements to traditional UAV sensors (SQ4). By drawing on biological principles and reviewing recent robotic implementations, this chapter established a foundational understanding of tactile sensing's role in aerial robotics. These insights set the stage for the next chapter, which investigates how neuromorphic computing, particularly SNNs, can be used to encode, process, and classify tactile signals in a power-efficient and real-time manner, further advancing the potential of autonomous aerial systems in challenging environments.

Spiking Neural Networks

This chapter provided an in-depth exploration of SNNs, emphasizing their biological foundations, computational advantages, and relevance for tactile perception in aerial robotics. SNNs differ fundamentally from ANNs by using discrete spike events rather than continuous activations to encode and process information (SQ2). This event-driven mechanism closely mimics the operation of biological neurons, enabling sparse, asynchronous activity that translates to low-latency and energy-efficient computation. Such characteristics make SNNs particularly suitable for embedded systems with strict power and real-time constraints.

The chapter examined the core components of neuromorphic pipelines, including input and output encoding strategies, training methodologies, and deployment on neuromorphic hardware. Tactile data from high-dimensional, time-varying sensors, particularly whisker-inspired arrays, can be effectively encoded into spike trains through methods such as rate coding, latency coding, and delta modulation (SQ5). Output decoding typically involves rate, latency, or population coding. To overcome the challenge of non-differentiable spikes during learning, several training strategies have been introduced, including shadow training, BPTT, biologically inspired local learning rules, and surrogate gradient methods. These approaches enable SNNs to learn temporal patterns and achieve robustness to noise while operating under low power budgets.

Through a review of recent applications, such as event-driven object recognition and tactile classification, the chapter demonstrated that SNNs are not only theoretically grounded but also increasingly viable in real-world robotic systems. When combined with tactile inputs, especially from whisker-inspired sensors, SNNs offer a biologically plausible and computationally efficient path forward for autonomous aerial platforms. While challenges remain in training scalability and hardware integration, the findings of this chapter underscore the potential of SNNs to serve as the processing backbone for next-generation, situation-aware UAVs.

5.2. Gaps in Literature

Critically reviewing the findings of the literature reseach leads us to gaps in the literature. In this section we present and discuss several important gaps Although significant progress has been made in tactile sensing and neuromorphic computing, several important gaps remain in the literature:

- **Lack of Whisker-Equipped UAV Platforms**

  Despite their promising potential, the integration of whisker-inspired sensors on UAV platforms remains significantly underexplored. Numerous studies have demonstrated the effectiveness of such sensors in ground-based robotic systems, including mobile robots, manipulators, and subterranean exploration platforms [82, 45, 39, 57, 56, 55]. Similar implementations have also emerged in underwater robotics [46, 21], yet their deployment in aerial systems is still rare. Existing studies addressing flow detection for UAVs typically evaluate sensor arrays in isolation rather than on-board drones [34]. Other related works focus on alternative tactile approaches, such as compliant robotic fingers [8] or bio-inspired airflow disturbance sensors [71], without employing whisker-like structures. Only one study has explored the use of haptic antennae for UAV navigation in confined spaces [53], highlighting the novelty of tactile aerial systems. Although biomorphic whisker designs specifically intended for aerial platforms have recently been proposed [79], they have not yet been validated through practical UAV experiments. This reveals a clear research gap in the development, integration, and evaluation of whisker-equipped UAVs capable of performing close-proximity tactile interactions in unstructured or low-visibility environments.

- **Limited Integration of SNNs with Whisker-Based Sensing**

  While SNNs and whisker-based tactile sensing have both been studied extensively in isolation, their combined application in robotic systems remains largely unexplored. SNNs have shown promising results in object detection tasks, particularly when paired with event-based visual or general tactile data [25, 30, 2]. Similarly, whisker-inspired sensors have been successfully employed for object recognition and exploration in various robotic platforms [62, 41, 77]. However, the direct integration of SNNs with robotic whisker arrays for tactile data processing is scarcely addressed in the literature. Existing work on SNNs with tactile sensors involves fingertip sensors or piezoresistive arrays [72, 11], rather than whisker-like modalities. While some studies have applied SNNs to process biological whisker data for understanding neural mechanisms in

animals [6, 58], no published research has demonstrated a fully integrated robotic system that combines artificial whisker sensors with SNN-based processing. This represents a significant gap and an opportunity for future work to develop energy-efficient, biologically plausible tactile systems capable of real-time perception in complex environments.

- **Absence of Fully Neuromorphic Tactile UAV Systems**

  To date, only a single study has demonstrated the integration of SNNs with neuromorphic sensors and hardware on a UAV platform. Paredes-Vallés et al. presented a fully neuromorphic vision and control framework enabling autonomous drone flight [54]. Beyond this example, there are no studies that combine neuromorphic processing with tactile or whisker-inspired sensing on UAVs. Consequently, the development of fully neuromorphic whiskered drones, integrating event-driven sensors, SNN-based processing, and neuromorphic hardware, remains an open and promising research direction.

5.3. Future Research Directions

This literature review explored the intersection of aerial tactile perception and neuromorphic computing, with particular emphasis on the potential of whisker-inspired sensors. The guiding research question was:

> *What are the emerging research directions in aerial tactile perception using whisker-inspired sensors and Spiking Neural Networks?*

To advance the field, future research should prioritize the development and experimental validation of neuromorphic whisker-equipped UAVs. Potential application domains include airflow sensing, aerial tactile exploration, and close-proximity mapping, which can all support autonomous navigation in cluttered, confined, or vision-denied environments. These scenarios would particularly benefit from the lightweight, low-power, and high-resolution processing capabilities offered by SNNs.

Although whisker-inspired tactile sensors and SNNs have each shown individual promise, their integration in aerial robotics remains largely unaddressed. The fusion of whisker arrays with neuromorphic computing for real-time, event-driven tactile perception represents a compelling and underdeveloped area of research. Such systems could deliver biologically plausible and energy-efficient sensing pipelines that enable UAVs to operate effectively in GNSS-denied and low-visibility environments. This holds for both flow-sensitive whiskers, which respond to air currents, and contact-sensitive whiskers, which detect physical interaction with surfaces.

Whisker-inspired flow sensors, with their sensitivity to multi-directional airflow and ability to distinguish between self-induced and external stimuli, offer unique advantages for enhancing UAV situational awareness, flight stability, and collision avoidance. Meanwhile, contact whiskers could enable high-fidelity tactile interaction. A neuromorphic tactile system, for example, could detect a whisker contact and immediately trigger evasive maneuvers, such as stopping or turning, to avoid obstacles in real time.

Despite this conceptual promise, several technical challenges remain. These include improving large-scale training methods for SNNs, closing the performance gap with state-of-the-art ANNs, and developing end-to-end neuromorphic pipelines that process raw event-based data without requiring frame-based conversions. Addressing these challenges is essential to realize responsive, efficient, and scalable tactile systems for UAV platforms.

Ultimately, the long-term vision is to develop animal-like tactile perception on UAVs by emulating the event-driven and parallel processing principles of biological systems. Such platforms could achieve fast, low-latency, and autonomous operation in unstructured environments. This research direction offers a promising pathway toward biologically inspired UAVs capable of navigating their environments with enhanced tactile awareness and minimal energy usage.

6

Research Proposal

While several future research directions mentioned in Chapter 5 are conceptually promising, some remain distant from practical implementation. In particular, the development of a fully neuromorphic tactile UAV capable of autonomous operation in unstructured environments is a long-term objective. A more feasible and immediate step is to investigate the integration of SNNs with whisker-inspired sensors, an area that remains largely unexplored yet holds strong potential for enabling energy-efficient, real-time tactile perception.

Whisker-based tactile perception represents a foundational capability for achieving robust and autonomous UAV navigation in unstructured or GNSS-denied environments, where conventional vision-based systems may fail. A natural progression from the identified gaps is to explore tactile object detection through integration of SNNs with a biomorphic whisker array. By focusing on object classification tasks, this study aims to provide a practical proof of concept for integrating whisker-based tactile sensing and neuromorphic processing. The primary objective of the proposed research is to develop and evaluate a SNN model that analyses tactile data from a biomorphic whisker array. The proposed system will classify objects relevant to UAV navigation and assess the feasibility of real-time, low-power tactile perception using SNNs.

This chapter is structured as follows: First, we discuss the scope and feasibility of the project, including technical considerations and constraints. Next, we present the system scenario that motivates and contextualizes the proposed research. Finally, we formulate the main research question that will guide the thesis project.

6.1. Scope and Feasibility

This study aims to develop a proof-of-concept system that combines a biomorphic whisker sensor array with a SNN to perform tactile object classification. The experiment will be conducted in a controlled laboratory environment using a fixed setup, in which objects are swept past a static array of whiskers mounted on a linear stage. At this stage, no deployment on a UAV or real-time inference on neuromorphic hardware (e.g., Intel Loihi) is required. The outcomes will serve as a foundational step toward future research on neuromorphic tactile UAVs .

The tactile dataset will be generated by sweeping three rigid object shapes (e.g., wall, sphere, and box) across a 10-whisker sensor array mounted on a rail-based linear actuator. The actuator will move the objects at a controlled speed, enabling reproducible contact events. Each shape is expected to produce a distinct spatio-temporal pressure pattern across the whisker array.

Sensor signals will be preprocessed using simple filtering techniques, including optional bandpass filtering to reduce noise and temperature compensation for sensor drift. The processed signals will be encoded into spike trains using rate encoding, although direct input of continuous signals into the SNN will also be explored for comparison.

A simple Multi-Layer Perceptron (MLP) model with LIF neurons will serve as the initial SNN architecture. The network will be trained with surrogate gradient descent and must classify the three object shapes based on the encoded tactile input. The output will use rate decoding, with the neuron exhibiting the highest spike count representing the predicted class.

Scenario	Tactile Perception with a Whisker-Inspired Sensor Array Using a Spiking Neural Network
System	A biomorphic whisker sensor array, combined with an MLP-SNN and GPU-based local training for tactile object detection.
Objectives	• Convert analog whisker signals into spike trains using rate coding • Create a tactile event-based dataset by recording whisker array interactions with various shapes. • Train an SNN to process tactile events for shape detection • Design and fabricate test objects for experiments • Evaluate system performance with accuracy metrics
Limitations	• Proof-of-concept only, simplified research • Offline computing; no neuromorphic chip available
Constraints	• Computational complexity of SNN training • Whisker sensor signal issues (e.g., hysteresis, noise)
Concepts	• Whisker-inspired sensor array • Spiking Neural Network • Tactile object detection

Table 6.1: Proposed system scenario.

Model performance will be assessed using accuracy metrics on a train/test split of the dataset, along with confusion matrices to visualize misclassifications. Additional evaluations will include robustness checks, ablation studies, and hyperparameter searches to optimize performance.

6.2. System Scenario

The system scenario, including its objectives and constraints, is presented in Table 6.1.

6.3. Research Question

The following research question has been proposed for the MSc thesis project:

> *"To what extent can a MLP-SNN model classify a wall, an inclined wall and a spherical object based on pressure signals from a 10-whisker biomorphic sensor array? "*

We will quantify performance using: (i) rate-coding output decoding accuracy, (ii) accuracy on train/validation/test sets and (iii) confusion matrices. To assess the contribution of the SNN component, we will compare against a non-spiking MLP with the same feature inputs. Architectural ablations (e.g., MLP-only vs. hybrid MLP–SNN) will be included. The dataset comprises three classes: sphere, wall (planar), and inclined wall (planar with fixed tilt). Each sample is a multivariate whisker-pressure time series from the 10-sensor array, labeled by ground truth object class. If the proof-of-concept classification is feasible, we will extend toward tactile mapping: from discrete classification to regression-based surface reconstruction. Additional directions include alternative whisker configurations, richer input features, and other SNN architectures.

This would form a critical first step toward enabling autonomous UAV navigation in unstructured environments. A new world opens up with the advent of extremely-fast and energy-efficient computing on robots. This advancement enables perception to approach the speed of natural sensory systems, allowing the development of animal-like robots with advanced perceptual capabilities. Such systems can operate in a biologically plausible and energy-efficient manner, pushing the boundaries of robotics.

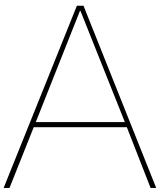

Use of Generative AI

In composing this literature assignment, I employed OpenAI's language model, ChatGPT, as a tool to improve the linguistic quality of the document. This AI-powered tool played a key role in enhancing sentence structures, expanding vocabulary, and ensuring grammatical precision, thus boosting the overall readability and professional quality of the text. It is important to note, however, that while ChatGPT assisted in refining the clarity and style of the text, the concepts and arguments presented are entirely my own.

Bibliography

[1] A. S. Ahl. “The role of vibrissae in behavior: A status review”. In: *Veterinary Research Communications* 10.1 (Dec. 1, 1986), pp. 245–268. ISSN: 1573-7446. DOI: 10.1007/BF02213989. URL: https://doi.org/10.1007/BF02213989 (visited on 05/20/2025).

[2] Soikat Hasan Ahmed, Jan Finkbeiner, and Emre Neftci. *Efficient Event-Based Object Detection: A Hybrid Neural Network with Spatial and Temporal Attention*. Mar. 11, 2025. DOI: 10.48550/arXiv.2403.10173. arXiv: 2403.10173[cs]. URL: http://arxiv.org/abs/2403.10173 (visited on 06/05/2025).

[3] Filipp Akopyan et al. “TrueNorth: Design and Tool Flow of a 65 mW 1 Million Neuron Programmable Neurosynaptic Chip”. In: *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 34.10 (Oct. 2015), pp. 1537–1557. ISSN: 1937-4151. DOI: 10.1109/TCAD.2015.2474396. URL: https://ieeexplore.ieee.org/abstract/document/7229264 (visited on 07/14/2025).

[4] Steffen Albrecht et al. *Spiking neural networks provide accurate, efficient and robust models for whisker stimulus classification and allow for inter-individual generalization*. en. Pages: 2023.04.19.537473 Section: New Results. Apr. 2023. DOI: 10.1101/2023.04.19.537473. URL: https://www.biorxiv.org/content/10.1101/2023.04.19.537473v1 (visited on 05/29/2025).

[5] Leiv Andresen, Emanuele Aucone, and Stefano Mintchev. “Whisker-based Haptic Perception System for Branch Detection in Dense Vegetation”. In: *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*. 2022 IEEE 5th International Conference on Soft Robotics (RoboSoft). Apr. 2022, pp. 911–918. DOI: 10.1109/RoboSoft54090.2022.9762143. URL: https://ieeexplore.ieee.org/abstract/document/9762143 (visited on 06/13/2025).

[6] Alberto Antonietti et al. “Brain-Inspired Spiking Neural Network Controller for a Neurorobotic Whisker System”. In: *Frontiers in Neurorobotics* 16 (June 13, 2022). Publisher: Frontiers. ISSN: 1662-5218. DOI: 10.3389/fnbot.2022.817948. URL: https://www.frontiersin.org/journals/neurorobotics/articles/10.3389/fnbot.2022.817948/full (visited on 05/20/2025).

[7] George Brayshaw, Benjamin Ward-Cherrier, and Martin J. Pearson. “A Neuromorphic System for the Real-time Classification of Natural Textures”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. 2024 IEEE International Conference on Robotics and Automation (ICRA). May 2024, pp. 1070–1076. DOI: 10.1109/ICRA57147.2024.10610401. URL: https://ieeexplore.ieee.org/document/10610401 (visited on 06/05/2025).

[8] Anton Bredenbeck, Cosimo Della Santina, and Salua Hamaza. “Embodying Compliant Touch on Drones for Aerial Tactile Navigation”. In: *IEEE Robotics and Automation Letters* 10.2 (Feb. 2025), pp. 1209–1216. ISSN: 2377-3766. DOI: 10.1109/LRA.2024.3519888. URL: https://ieeexplore-ieee-org.tudelft.idm.oclc.org/document/10806756 (visited on 05/19/2025).

[9] Natalia Caporale and Yang Dan. “Spike Timing–Dependent Plasticity: A Hebbian Learning Rule”. In: *Annual Review of Neuroscience* 31.1 (July 1, 2008), pp. 25–46. ISSN: 0147-006X, 1545-4126. DOI: 10.1146/annurev.neuro.31.060407.125639. URL: https://www.annualreviews.org/doi/10.1146/annurev.neuro.31.060407.125639 (visited on 10/07/2025).

[10] CNN. *Space Robots: The Future of Exploration*. Accessed: 2025-07-11. 2013. URL: https://edition.cnn.com/2013/12/23/tech/space-robots-future.

[11] Ali Dabbous et al. “Object Contact Shape Classification Using Neuromorphic Spiking Neural Network with STDP Learning”. In: *2022 IEEE International Symposium on Circuits and Systems (ISCAS)*. ISSN: 2158-1525. May 2022, pp. 1052–1056. DOI: 10.1109/ISCAS48785.2022.9937733. URL: https://ieeexplore.ieee.org/abstract/document/9937733 (visited on 06/18/2025).

[12] Mike Davies et al. “Advancing Neuromorphic Computing With Loihi: A Survey of Results and Outlook”. In: *Proceedings of the IEEE* 109.5 (May 2021), pp. 911–934. ISSN: 1558-2256. DOI: 10.1109/JPROC.2021.3067593. URL: https://ieeexplore-ieee-org.tudelft.idm.oclc.org/abstract/document/9395703 (visited on 07/14/2025).

[13] William Deer and Pauline E. I. Pounds. “Lightweight Whiskers for Contact, Pre-Contact, and Fluid Velocity Sensing”. In: *IEEE Robotics and Automation Letters* 4.2 (Apr. 2019), pp. 1978–1984. ISSN: 2377-3766. DOI: 10.1109/LRA.2019.2899215. URL: https://ieeexplore.ieee.org/abstract/document/8641341 (visited on 05/29/2025).

[14] Mathew E. Diamond et al. “’Where’ and ’what’ in the whisker sensorimotor system | Nature Reviews Neuroscience”. In: (Aug. 2008). DOI: 10.1038/nrn2411. URL: https://www.nature.com/articles/nrn2411 (visited on 06/18/2025).

[15] Hannah Emnett, Matthew Graff, and Mitra Hartmann. “A Novel Whisker Sensor Used for 3D Contact Point Determination and Contour Extraction”. In: *Robotics Science and Systems* XIV (June 2018 June 2018). URL: https://par.nsf.gov/biblio/10073510-novel-whisker-sensor-used-contact-point-determination-contour-extraction (visited on 06/02/2025).

[16] Jason Eshraghian et al. *snnTorch Documentation*. Accessed: 2025-09-24. snnTorch Project. 2023. URL: https://snntorch.readthedocs.io/en/latest/readme.html.

[17] Jason K. Eshraghian et al. “Training Spiking Neural Networks Using Lessons From Deep Learning”. In: *Proceedings of the IEEE* 111.9 (Sept. 2023), pp. 1016–1054. ISSN: 1558-2256. DOI: 10.1109/JPROC.2023.3308088. URL: https://ieeexplore.ieee.org/abstract/document/10242251 (visited on 04/24/2025).

[18] Jeremy A. Fishel and Gerald 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)*. 2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob). ISSN: 2155-1782. June 2012, pp. 1122–1127. DOI: 10.1109/BioRob.2012.6290741. URL: https://ieeexplore.ieee.org/abstract/document/6290741 (visited on 07/03/2025).

[19] Charles Fox et al. “Tactile SLAM with a biomimetic whiskered robot”. In: *2012 IEEE International Conference on Robotics and Automation*. 2012 IEEE International Conference on Robotics and Automation. ISSN: 1050-4729. May 2012, pp. 4925–4930. DOI: 10.1109/ICRA.2012.6224813. URL: https://ieeexplore.ieee.org/abstract/document/6224813 (visited on 07/02/2025).

[20] Fangming Guo et al. “Event-Driven Tactile Sensing With Dense Spiking Graph Neural Networks”. In: *IEEE Transactions on Instrumentation and Measurement* 74 (2025), pp. 1–13. ISSN: 1557-9662. DOI: 10.1109/TIM.2025.3541787. URL: https://ieeexplore.ieee.org/abstract/document/10884798 (visited on 06/18/2025).

[21] Linan Guo et al. “Piezoelectric wavy whisker sensor for perceiving underwater vortex from a bluff body”. In: *Sensors and Actuators A: Physical* 365 (Jan. 1, 2024), p. 114875. ISSN: 0924-4247. DOI: 10.1016/j.sna.2023.114875. URL: https://www.sciencedirect.com/science/article/pii/S0924424723007240 (visited on 06/02/2025).

[22] Salua Hamaza, Ioannis Georgilas, and Thomas Richardson. “2D Contour Following with an Unmanned Aerial Manipulator: Towards Tactile-Based Aerial Navigation”. In: *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). ISSN: 2153-0866. Nov. 2019, pp. 3664–3669. DOI: 10.1109/IROS40897.2019.8968591. URL: https://ieeexplore.ieee.org/abstract/document/8968591 (visited on 06/24/2025).

[23] A. L. Hodgkin and A. F. Huxley. “A quantitative description of membrane current and its application to conduction and excitation in nerve”. In: *The Journal of Physiology* 117.4 (Aug. 28, 1952), pp. 500–544. ISSN: 0022-3751. DOI: 10.1113/jphysiol.1952.sp004764. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1392413/ (visited on 07/10/2025).

[24] Alex C. Hollenbeck et al. “Bioinspired Artificial Hair Sensors for Flight-by-Feel of Unmanned Aerial Vehicles: A Review”. In: (Dec. 2023). DOI: 10.2514/1.J062931. URL: https://arc.aiaa.org/doi/epdf/10.2514/1.J062931 (visited on 06/02/2025).

[25] Jiarui Hu et al. “Tacformer: A Self-attention Spiking Neural Network for Tactile Object Recognition”. In: *Intelligent Robotics and Applications*. International Conference on Intelligent Robotics and Applications. ISSN: 1611-3349. Springer, Singapore, 2023, pp. 156–168. ISBN: 978-981-99-6495-6. DOI: 10.1007/978-981-99-6495-6_14. URL: https://link.springer.com/chapter/10.1007/978-981-99-6495-6_14 (visited on 06/18/2025).

[26] Lucie A. Huet, John W. Rudnicki, and Mitra J. Z. Hartmann. “Tactile Sensing with Whiskers of Various Shapes: Determining the Three-Dimensional Location of Object Contact Based on Mechanical Signals at the Whisker Base | Soft Robotics”. In: (2017). DOI: 10.1089/soro.2016.0028. URL: https://www-liebertpub-com.tudelft.idm.oclc.org/doi/abs/10.1089/soro.2016.0028 (visited on 06/02/2025).

[27] H. Hyvärinen et al. “Aquatic Environment and Differentiation of Vibrissae: Comparison of Sinus Hair Systems of Ringed Seal, Otter and Pole Cat”. In: *Brain Behavior and Evolution* 74.4 (Dec. 8, 2009), pp. 268–279. ISSN: 0006-8977. DOI: 10.1159/000264662. URL: https://doi.org/10.1159/000264662 (visited on 09/29/2025).

[28] E.M. Izhikevich. “Simple model of spiking neurons”. In: *IEEE Transactions on Neural Networks* 14.6 (Nov. 2003), pp. 1569–1572. ISSN: 1941-0093. DOI: 10.1109/TNN.2003.820440. URL: https://ieeexplore.ieee.org/abstract/document/1257420 (visited on 07/14/2025).

[29] Nawid Jamali and Claude Sammut. “Material classification by tactile sensing using surface textures”. In: *2010 IEEE International Conference on Robotics and Automation*. 2010 IEEE International Conference on Robotics and Automation. ISSN: 1050-4729. May 2010, pp. 2336–2341. DOI: 10.1109/ROBOT.2010.5509675. URL: https://ieeexplore.ieee.org/abstract/document/5509675 (visited on 07/03/2025).

[30] Xiaobo Jin et al. “R-SNN: Region-Based Spiking Neural Network for Object Detection”. In: *IEEE Transactions on Cognitive and Developmental Systems* 16.3 (June 2024), pp. 810–817. ISSN: 2379-8939. DOI: 10.1109/TCDS.2023.3311634. URL: https://ieeexplore.ieee.org/abstract/document/10244119 (visited on 06/05/2025).

[31] D. Jung and A. Zelinsky. “Whisker based mobile robot navigation”. In: *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS ’96*. IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS ’96. Vol. 2. Nov. 1996, 497–504 vol.2. DOI: 10.1109/IROS.1996.570842. URL: https://ieeexplore.ieee.org/abstract/document/570842 (visited on 05/20/2025).

[32] Teresa Kent. “Whisker-Inspired Sensors for Unstructured Environments”. thesis. Carnegie Mellon University, Mar. 7, 2025. DOI: 10.1184/R1/28513922.v1. URL: https://kilthub.cmu.edu/articles/thesis/Whisker-Inspired_Sensors_for_Unstructured_Environments/28513922/1 (visited on 06/13/2025).

[33] Teresa A Kent et al. “WhiskSight: A Reconfigurable, Vision-Based, Optical Whisker Sensing Array for Simultaneous Contact, Airflow, and Inertia Stimulus Detection”. In: *IEEE Robotics and Automation Letters* 6.2 (Apr. 2021), pp. 3357–3364. ISSN: 2377-3766. DOI: 10.1109/LRA.2021.3062816. URL: https://ieeexplore.ieee.org/abstract/document/9366394 (visited on 06/02/2025).

[34] Teresa A. Kent and Sarah Bergbreiter. “Flow Shadowing: A Method to Detect Multiple Flow Headings using an Array of Densely Packed Whisker-inspired Sensors”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. 2024 IEEE International Conference on Robotics and Automation (ICRA). May 2024, pp. 17843–17849. DOI: 10.1109/ICRA57147.2024.10610885. URL: https://ieeexplore.ieee.org/abstract/document/10610885 (visited on 06/02/2025).

[35] DaeEun Kim and Ralf Möller. “Biomimetic whiskers for shape recognition”. In: *Robotics and Autonomous Systems* 55.3 (Mar. 31, 2007), pp. 229–243. ISSN: 0921-8890. DOI: 10.1016/j.robot.2006.08.001. URL: https://www.sciencedirect.com/science/article/pii/S0921889006001400 (visited on 06/02/2025).

[36] Jaehun Kim et al. “Object shape recognition using tactile sensor arrays by a spiking neural network with unsupervised learning”. In: *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. 2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC). ISSN: 2577-1655. Oct. 2020, pp. 178–183. DOI: 10.1109/SMC42975.2020.9283337. URL: https://ieeexplore.ieee.org/abstract/document/9283337 (visited on 06/18/2025).

[37] Suhan Kim et al. “A Whisker-inspired Fin Sensor for Multi-directional Airflow Sensing”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). ISSN: 2153-0866. Oct. 2020, pp. 1330–1337. DOI: 10.1109/IROS45743.2020.9341723. URL: https://ieeexplore.ieee.org/abstract/document/9341723 (visited on 05/20/2025).

[38] Taekyoung Kim et al. “Soft Airflow Sensors With Artificial Hair Structures and Printed Ionogel Channels for Wind Gust Detection for Small Uncrewed Vehicles”. In: *IEEE/ASME Transactions on Mechatronics* 28.4 (Aug. 2023), pp. 2317–2327. ISSN: 1941-014X. DOI: 10.1109/TMECH.2022.3233669. URL: https://ieeexplore.ieee.org/abstract/document/10013774 (visited on 06/02/2025).

[39] Tanel Kossas et al. “Whisker-Based Tactile Navigation Algorithm For Underground Robots”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. 2024 IEEE International Conference on Robotics and Automation (ICRA). May 2024, pp. 13164–13170. DOI: 10.1109/ICRA57147.2024.10610762. URL: https://ieeexplore.ieee.org/document/10610762/ (visited on 06/30/2025).

[40] Gauri Kshettry et al. “TuneNav: Tunable Whisker Array for Touch-based Navigation in Confined Spaces”. In: *2024 IEEE MIT Undergraduate Research Technology Conference (URTC)*. 2024 IEEE MIT Undergraduate Research Technology Conference (URTC). Oct. 2024, pp. 1–5. DOI: 10.1109/URTC65039.2024.10937626. URL: https://ieeexplore.ieee.org/abstract/document/10937626 (visited on 06/30/2025).

[41] Alexander Kugele et al. “Hybrid SNN-ANN: Energy-Efficient Classification and Object Detection for Event-Based Vision”. In: *Pattern Recognition*. Ed. by Christian Bauckhage, Juergen Gall, and Alexander Schwing. Cham: Springer International Publishing, 2021, pp. 297–312. ISBN: 978-3-030-92659-5. DOI: 10.1007/978-3-030-92659-5_19.

[42] Deepesh Kumar, Andrei Nakagawa Silva, and Nitish V. Thakor. “Neuromorphic Tactile Sensing and Encoding”. In: *Handbook of Neuroengineering*. Springer, Singapore, 2023, pp. 1609–1634. ISBN: 978-981-16-5540-1. DOI: 10.1007/978-981-16-5540-1_117. URL: https://link.springer.com/rwe/10.1007/978-981-16-5540-1_117 (visited on 07/14/2025).

[43] Nathan F. Lepora, Martin Pearson, and Luke Cramphorn. “TacWhiskers: Biomimetic Optical Tactile Whiskered Robots”. In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). ISSN: 2153-0866. Oct. 2018, pp. 7628–7634. DOI: 10.1109/IROS.2018.8593653. URL: https://ieeexplore.ieee.org/abstract/document/8593653 (visited on 06/02/2025).

[44] Michael A. Lin et al. *Navigation and 3D Surface Reconstruction from Passive Whisker Sensing*. June 10, 2024. DOI: 10.48550/arXiv.2406.06038. arXiv: 2406.06038[cs]. URL: http://arxiv.org/abs/2406.06038 (visited on 06/30/2025).

[45] Michael A. Lin et al. “Whisker-Inspired Tactile Sensing for Contact Localization on Robot Manipulators”. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). ISSN: 2153-0866. Oct. 2022, pp. 7817–7824. DOI: 10.1109/IROS47612.2022.9982122. URL: https://ieeexplore.ieee.org/abstract/document/9982122 (visited on 05/20/2025).

[46] Jianhua Liu et al. “Whisker-inspired and self-powered triboelectric sensor for underwater obstacle detection and collision avoidance”. In: *Nano Energy* 101 (Oct. 1, 2022), p. 107633. ISSN: 2211-2855. DOI: 10.1016/j.nanoen.2022.107633. URL: https://www.sciencedirect.com/science/article/pii/S221128552200711X (visited on 06/19/2025).

[47] Kai Malcolm and Josue Casco-Rodriguez. *A Comprehensive Review of Spiking Neural Networks: Interpretation, Optimization, Efficiency, and Best Practices*. Mar. 21, 2023. DOI: 10.48550/arXiv.2303.10780. arXiv: 2303.10780[cs]. URL: http://arxiv.org/abs/2303.10780 (visited on 04/24/2025).

[48] Michael J. Milford and Gordon F. Wyeth. “Mapping a Suburb With a Single Camera Using a Biologically Inspired SLAM System”. In: *IEEE Transactions on Robotics* 24.5 (Oct. 2008), pp. 1038–1053. ISSN: 1941-0468. DOI: 10.1109/TRO.2008.2004520. URL: https://ieeexplore.ieee.org/abstract/document/4627450 (visited on 07/09/2025).

[49] Ben Mitchinson et al. “Perception of Simple Stimuli Using Sparse Data from a Tactile Whisker Array”. In: *Biomimetic and Biohybrid Systems*. Conference on Biomimetic and Biohybrid Systems. ISSN: 1611-3349. Springer, Berlin, Heidelberg, 2013, pp. 179–190. ISBN: 978-3-642-39802-5. DOI: 10.1007/978-3-642-39802-5_16. URL: https://link.springer.com/chapter/10.1007/978-3-642-39802-5_16 (visited on 06/17/2025).

[50] Ravinder S. Mohan and Tony J. Prescott. “Biomimetic Whiskered Robots for Tactile Sensing”. In: *Sensors* 22.7 (2022), p. 2705. DOI: 10.3390/s22072705. URL: https://www.mdpi.com/1424-8220/22/7/2705.

[51] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. “Surrogate Gradient Learning in Spiking Neural Networks: Bringing the Power of Gradient-Based Optimization to Spiking Neural Networks”. In: *IEEE Signal Processing Magazine* 36.6 (Nov. 2019), pp. 51–63. ISSN: 1558-0792. DOI: 10.1109/MSP.2019.2931595. URL: https://ieeexplore.ieee.org/document/8891809 (visited on 05/19/2025).

[52] Eustace Painkras et al. “SpiNNaker: A 1-W 18-Core System-on-Chip for Massively-Parallel Neural Network Simulation”. In: *IEEE Journal of Solid-State Circuits* 48.8 (Aug. 2013), pp. 1943–1953. ISSN: 1558-173X. DOI: 10.1109/JSSC.2013.2259038. URL: https://ieeexplore.ieee.org/abstract/document/6515159 (visited on 07/14/2025).

[53] Christos Papachristos, Shehryar Khattak, and Kostas Alexis. “Haptic Feedback-Based Reactive Navigation for Aerial Robots Subject to Localization Failure”. In: *2019 IEEE Aerospace Conference*. 2019 IEEE Aerospace Conference. ISSN: 1095-323X. Mar. 2019, pp. 1–7. DOI: 10.1109/AERO.2019.8741634. URL: https://ieeexplore.ieee.org/abstract/document/8741634 (visited on 06/20/2025).

[54] F. Paredes-Vallés et al. “Fully neuromorphic vision and control for autonomous drone flight”. In: *Science Robotics* 9.90 (May 15, 2024). Publisher: American Association for the Advancement of Science, eadi0591. DOI: 10.1126/scirobotics.adi0591. URL: https://www.science.org/doi/full/10.1126/scirobotics.adi0591 (visited on 05/06/2025).

[55] Martin J. Pearson et al. “Biomimetic vibrissal sensing for robots”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 366.1581 (Nov. 12, 2011). Publisher: Royal Society, pp. 3085–3096. DOI: 10.1098/rstb.2011.0164. URL: https://royalsocietypublishing.org/doi/full/10.1098/rstb.2011.0164 (visited on 06/18/2025).

[56] Martin J. Pearson et al. “SCRATCHbot: Active Tactile Sensing in a Whiskered Mobile Robot”. In: *From Animals to Animats 11*. International Conference on Simulation of Adaptive Behavior. ISSN: 1611-3349. Springer, Berlin, Heidelberg, 2010, pp. 93–103. ISBN: 978-3-642-15193-4. DOI: 10.1007/978-3-642-15193-4_9. URL: https://link.springer.com/chapter/10.1007/978-3-642-15193-4_9 (visited on 06/17/2025).

[57] Martin J. Pearson et al. “Whiskerbot: A Robotic Active Touch System Modeled on the Rat Whisker Sensory System”. In: *Adaptive Behavior* 15.3 (Sept. 1, 2007). Publisher: SAGE Publications Ltd STM, pp. 223–240. ISSN: 1059-7123. DOI: 10.1177/1059712307082089. URL: https://doi.org/10.1177/1059712307082089 (visited on 06/17/2025).

[58] Horst Petschenig et al. “Classification of Whisker Deflections From Evoked Responses in the Somatosensory Barrel Cortex With Spiking Neural Networks”. In: *Frontiers in Neuroscience* 16 (Apr. 14, 2022). Publisher: Frontiers. ISSN: 1662-453X. DOI: 10.3389/fnins.2022.838054. URL: https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2022.838054/full (visited on 05/20/2025).

[59] Tony J. Prescott, Paul F. Verschure, and Tony Belpaeme. “Neurorobotics: A Thriving Community and a Promising Pathway Toward Intelligent Cognitive Robots”. In: *IEEE Robotics & Automation Magazine* 25.3 (2018), pp. 9–16. DOI: 10.1109/MRA.2018.2849605. URL: https://www.researchgate.net/publication/326418438.

[60] Tony J. Prescott et al. “Whisking with robots”. In: *IEEE Robotics & Automation Magazine* 16.3 (Sept. 2009), pp. 42–50. ISSN: 1558-223X. DOI: 10.1109/MRA.2009.933624. URL: https://ieeexplore.ieee.org/abstract/document/5233415 (visited on 06/02/2025).

[61] Brian W. Quist and Mitra J. Z. Hartmann. “Mechanical signals at the base of a rat vibrissa: the effect of intrinsic vibrissa curvature and implications for tactile exploration”. In: *Journal of Neurophysiology* 107.9 (May 2012). Publisher: American Physiological Society, pp. 2298–2312. ISSN: 0022-3077. DOI: 10.1152/jn.00372.2011. URL: https://journals.physiology.org/doi/full/10.1152/jn.00372.2011 (visited on 09/29/2025).

[62] R. Andrew Russell and Jaury Adi Wijaya. “Recognising and manipulating objects using data from a whisker sensor array”. In: *Robotica* 23.5 (Sept. 2005), pp. 653–664. ISSN: 1469-8668, 0263-5747. DOI: 10.1017/S0263574704000748. URL: https://www.cambridge.org/core/journals/robotica/article/recognising-and-manipulating-objects-using-data-from-a-whisker-sensor-array/121F1D0AB2031AEE4E845642F96982D7 (visited on 06/13/2025).

[63] R.A. Russell. “Using tactile whiskers to measure surface contours”. In: *Proceedings 1992 IEEE International Conference on Robotics and Automation*. 1992 IEEE International Conference on Robotics and Automation. May 1992, 1295–1299 vol.2. DOI: 10.1109/ROBOT.1992.220070. URL: https://ieeexplore.ieee.org/document/220070 (visited on 06/19/2025).

[64] Mohammed Salman and Martin J. Pearson. “Advancing whisker based navigation through the implementation of Bio-Inspired whisking strategies”. In: *2016 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. 2016 IEEE International Conference on Robotics and Biomimetics (ROBIO). Dec. 2016, pp. 767–773. DOI: 10.1109/ROBIO.2016.7866416. URL: https://ieeexplore.ieee.org/document/7866416/ (visited on 06/30/2025).

[65] Hian Hian See et al. *ST-MNIST – The Spiking Tactile MNIST Neuromorphic Dataset*. May 8, 2020. DOI: 10.48550/arXiv.2005.04319. arXiv: 2005.04319[cs]. URL: http://arxiv.org/abs/2005.04319 (visited on 04/24/2025).

[66] Joseph H. Solomon and Mitra J. Hartmann. “Robotic whiskers used to sense features”. In: *Nature* 443.7111 (Oct. 2006). Publisher: Nature Publishing Group, pp. 525–525. ISSN: 1476-4687. DOI: 10.1038/443525a. URL: https://www.nature.com/articles/443525a (visited on 05/28/2025).

[67] Joseph H. Solomon and Mitra J. Z. Hartmann. “Extracting Object Contours with the Sweep of a Robotic Whisker Using Torque Information”. In: *The International Journal of Robotics Research* 29.9 (Aug. 1, 2010). Publisher: SAGE Publications Ltd STM, pp. 1233–1245. ISSN: 0278-3649. DOI: 10.1177/0278364908104468. URL: https://doi.org/10.1177/0278364908104468 (visited on 06/02/2025).

[68] Shiva Subbulakshmi Radhakrishnan et al. “A biomimetic neural encoder for spiking neural network”. In: *Nature Communications* 12.1 (Apr. 9, 2021). Publisher: Nature Publishing Group, p. 2143. ISSN: 2041-1723. DOI: 10.1038/s41467-021-22332-8. URL: https://www.nature.com/articles/s41467-021-22332-8 (visited on 06/17/2025).

[69] J. Charles Sullivan et al. “Tactile Discrimination Using Active Whisker Sensors”. In: *IEEE Sensors Journal* 12.2 (Feb. 2012), pp. 350–362. ISSN: 1558-1748. DOI: 10.1109/JSEN.2011.2148114. URL: https://ieeexplore.ieee.org/abstract/document/5756632 (visited on 07/03/2025).

[70] Sudharshan Suresh et al. “Tactile SLAM: Real-time inference of shape and pose from planar pushing”. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 2021 IEEE International Conference on Robotics and Automation (ICRA). ISSN: 2577-087X. May 2021, pp. 11322–11328. DOI: 10.1109/ICRA48506.2021.9562060. URL: https://ieeexplore.ieee.org/abstract/document/9562060 (visited on 05/20/2025).

[71] Andrea Tagliabue et al. “Touch the Wind: Simultaneous Airflow, Drag and Interaction Sensing on a Multirotor”. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). ISSN: 2153-0866. Oct. 2020, pp. 1645–1652. DOI: 10.1109/IROS45743.2020.9341797. URL: https://ieeexplore.ieee.org/abstract/document/9341797 (visited on 05/20/2025).

[72] Tasbolat Taunyazov et al. *Event-Driven Visual-Tactile Sensing and Learning for Robots*. arXiv:2009.07083 [cs]. Sept. 2020. DOI: 10.48550/arXiv.2009.07083. URL: http://arxiv.org/abs/2009.07083 (visited on 06/17/2025).

[73] TU Delft. *Flying Robots Survey Biodiversity and Climate Inside Tropical Rainforests*. Accessed: 2025-09-24. 2023. URL: https://www.tudelft.nl/en/2023/lr/flying-robots-survey-biodiversity-and-climate-inside-tropical-rainforests.

[74] Benjamin Ward-Cherrier, Nicholas Pestell, and Nathan F. Lepora. “NeuroTac: A Neuromorphic Optical Tactile Sensor applied to Texture Recognition”. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. 2020 IEEE International Conference on Robotics and Automation (ICRA). ISSN: 2577-087X. May 2020, pp. 2654–2660. DOI: 10.1109/ICRA40945.2020.9197046. URL: https://ieeexplore.ieee.org/abstract/document/9197046 (visited on 04/24/2025).

[75] Benjamin Ward-Cherrier et al. “The TacTip Family: Soft Optical Tactile Sensors with 3D-Printed Biomimetic Morphologies”. In: *Soft Robotics* 5.2 (Apr. 2018). Publisher: Mary Ann Liebert, Inc., publishers, pp. 216–227. ISSN: 2169-5172. DOI: 10.1089/soro.2017.0052. URL: https://www.liebertpub.com/doi/full/10.1089/soro.2017.0052 (visited on 07/03/2025).

[76] Jaury A. Wijaya and R. Andrew Russell. “Object exploration using whisker sensors”. In: (Nov. 2002). URL: https://scholar.google.com/scholar?hl=nl&as_sdt=0%2C5&q=Object+exploration+using+whisker+sensors&btnG=&inst=6173373803492361994 (visited on 06/02/2025).

[77] Chenxi Xiao et al. “Active Multiobject Exploration and Recognition via Tactile Whiskers”. In: *IEEE Transactions on Robotics* 38.6 (Dec. 2022), pp. 3479–3497. ISSN: 1941-0468. DOI: 10.1109/TRO.2022.3182487. URL: https://ieeexplore.ieee.org/abstract/document/9813357 (visited on 05/20/2025).

[78] Xingchen Xu, Nathan F. Lepora, and Benjamin Ward-Cherrier. “A Neuromorphic Tactile System for Reliable Braille Reading in Noisy Environments”. In: *IEEE Robotics and Automation Letters* 10.6 (June 2025), pp. 5225–5232. ISSN: 2377-3766. DOI: 10.1109/LRA.2025.3558707. URL: https://ieeexplore.ieee.org/abstract/document/10955194 (visited on 06/05/2025).

[79] Chaoxiang Ye, Guido De Croon, and Salua Hamaza. “A Biomorphic Whisker Sensor for Aerial Tactile Applications”. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. 2024 IEEE International Conference on Robotics and Automation (ICRA). May 2024, pp. 5257–5263. DOI: 10.1109/ICRA57147.2024.10610850. URL: https://ieeexplore.ieee.org/document/10610850 (visited on 04/30/2025).

[80] Zhengkun Yi, Yilei Zhang, and Jan Peters. “Biomimetic tactile sensors and signal processing with spike trains: A review”. In: *Sensors and Actuators A: Physical* 269 (Jan. 1, 2018), pp. 41–52. ISSN: 0924-4247. DOI: 10.1016/j.sna.2017.09.035. URL: https://www.sciencedirect.com/science/article/pii/S0924424717307689 (visited on 06/18/2025).

[81] Zhenhua Yu, Peter R. N. Childs, and Thrishantha Nanayakkara. *Towards the Neuromorphic Computing for Offroad Robot Environment Perception and Navigation*. May 5, 2023. DOI: 10.48550/arXiv.2305.03860. arXiv: 2305.03860[cs]. URL: http://arxiv.org/abs/2305.03860 (visited on 05/20/2025).

[82] Zhenhua Yu et al. “A Method to use Nonlinear Dynamics in a Whisker Sensor for Terrain Identification by Mobile Robots”. In: *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). ISSN: 2153-0866. Sept. 2021, pp. 8437–8443. DOI: 10.1109/IROS51168.2021.9636571. URL: https://ieeexplore.ieee.org/abstract/document/9636571 (visited on 06/02/2025).

[83] Zhenhua Yu et al. “Whisker Sensor for Robot Environments Perception: A Review”. In: *IEEE Sensors Journal* 24.18 (Sept. 2024), pp. 28504–28521. ISSN: 1558-1748. DOI: 10.1109/JSEN.2024.3429373. URL: https://ieeexplore.ieee.org/abstract/document/10602791 (visited on 06/02/2025).

[84] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. “GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force”. In: *Sensors* 17.12 (Dec. 2017). Number: 12 Publisher: Multidisciplinary Digital Publishing Institute, p. 2762. ISSN: 1424-8220. DOI: 10.3390/s17122762. URL: https://www.mdpi.com/1424-8220/17/12/2762 (visited on 07/14/2025).

[85] Chenlu Zhao, Qi Jiang, and Yibin Li. “A novel biomimetic whisker technology based on fiber Bragg grating and its application”. In: *Measurement Science and Technology* 28.9 (Aug. 2017). Publisher: IOP Publishing, p. 095104. ISSN: 0957-0233. DOI: 10.1088/1361-6501/aa7d36. URL: https://dx.doi.org/10.1088/1361-6501/aa7d36 (visited on 06/19/2025).

[86] Chenlin Zhou et al. “Direct training high-performance deep spiking neural networks: a review of theories and methods”. In: *Frontiers in Neuroscience* 18 (July 31, 2024). Publisher: Frontiers. ISSN: 1662-453X. DOI: 10.3389/fnins.2024.1383844. URL: https://www.frontiersin.orghttps://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2024.1383844/full (visited on 04/24/2025).

A

AI Statement

This thesis was completed with the support of several AI tools, used deliberately and selectively across different phases of the project. All ideas, interpretations, and conclusions presented in this thesis are my own.

The following tools were used:

Academic writing Claude Sonnet & Opus (Anthropic) — used to support drafting, editing, and improving the clarity of written sections.

Software assistance Claude Sonnet (Anthropic) — used to assist with code development, debugging, and implementation of the experimental pipeline.

Brainstorming and general knowledge ChatGPT (OpenAI) — used for exploratory discussion and background orientation during early stages of the project.