Università
di Genova

PHD PROGRAM IN SCIENCE AND TECHNOLOGY FOR ELECTRONIC
AND TELECOMMUNICATION ENGINEERING

Smart Human Machine Interface Using Piezoelectric Sensors and Artificial Intelligence

Hussein Bassal

Thesis submitted for the degree of *Doctor of Philosophy* (38° cycle)

14 May 2026

Prof. Maurizio VALLE	Tutor
Dr. Yahya Abbass	Co-tutor
Dr. Nicolò Boccardo	Co-tutor
Prof. Strahinja Dosen	Co-tutor
Prof. Maurizio VALLE	Head of the PhD program

Università
di Genova

DITEN DIPARTIMENTO
DI INGEGNERIA NAVALE, ELETTRICA,
ELETTRONICA E DELLE TELECOMUNICAZIONI

SMART HUMAN MACHINE INTERFACE USING PIEZOELECTRIC SENSORS AND ARTIFICIAL INTELLIGENCE

Hussein Bassal

⋈

Reviewers
Gian Domenico LICCIARDO, Full Professor, University of Salerno, Italy
Roberto MEATTINI, Assistant Professor, University of Bologna, Italy

Examiners

Salvatore Pirozzi, Full Professor, Università degli Studi della Campania – Luigi Vanvitelli, Italy
Paolo Motto Ros, Assistant Professor, Politecnico di Torino, Italy
Paolo Gastaldo, Full Professor, Università di Genova, Italy

University of Genoa

Hussein Bassal

Smart Human Machine Interface Using Piezoelectric Sensors and Artificial Intelligence

xix+188 p.

Declaration I declare that, except where explicit reference is made to the work of others, this dissertation is entirely my own original work and has not been submitted, in whole or in part, for any other degree or qualification at this or any other university. It contains no material produced in collaboration with others, except where such collaboration is clearly indicated in the text and in the Acknowledgements. The dissertation comprises fewer than 65,000 words (including the bibliography, footnotes, tables, and equations) and includes fewer than 150 figures.

Abstract

The sense of touch is essential for stable object manipulation, safe interaction, and embodiment. In upper-limb prosthetics, however, tactile sensation is largely absent, forcing users to rely heavily on vision and increasing cognitive load. Addressing this limitation requires not only tactile sensors, but also mechanically compliant and scalable hand sensorization, reliable real-time decoding of tactile signals under variable contact conditions, and feedback-ready outputs that can be translated into intuitive stimulation patterns.

This thesis develops and validates a tactile sensing framework for upper-limb prosthetic hands based on piezoelectric polymer (PVDF) sensing, progressing from localized biomimetic sensing to prosthetic-hand integration, and finally to a distributed tactile sensing and electrotactile feedback architecture that combines scalable hand sensorization, real-time signal-processing pipelines, and multichannel sensory stimulation. First, a biomimetic PVDF-based fingertip sensing system is used for hardness discrimination through machine learning, with particular attention to incremental decision-making. On a fixed-size window dataset, the best-performing model achieves 94.6% Top-1 accuracy on 11,869 test windows, while Top-2 accuracy increases to 98.2%, indicating that most errors are near-miss confusions.

Second, the overall approach has been transferred to the Hannes prosthetic hand, where a real-time PC-based pipeline is developed for signal acquisition, conditioning, event-driven segmentation, and hardness inference using a 1D convolutional neural network. In offline evaluation with non-overlapping 100 ms windows across five hardness levels, the CNN achieves 91.4% Top-1 accuracy on 895 test windows, increasing to 96.8% and 98.7% for Top-2 and Top-3, respectively. Reliability analysis shows that very high-confidence outputs are correct in 97.5% of cases. In online deployment, majority voting over windows yields 80.0% event-level accuracy over 15 grasps, while confidence-gated detection provides at least one correct high-confidence window in 80.0% of grasps with a median latency of 0.21 s, demonstrating sub-second evidence accumulation.

Third, the work advances to a distributed tactile sensing and feedback architecture that combines 64 sensors integrated across the fingertips and palm, a real-time processing pipeline for binary contact-event extraction, and a multichannel non-invasive electrotactile feedback interface. These events drive a 64-channel stimulator and four 4×4 matrix electrodes using a nearly one-to-one mapping between sensing units and stimulation pads, thereby preserving the topographic organization of tactile information. Validation in healthy subjects ($n = 5$) across five perceptual tasks shows strong spatial performance for finger/palm detection (94.7±6.5%) and robust discrimination of grasp-related contact extent and closure (82.0±5.7% and 82.7±10.1%), whereas tasks dominated by temporal cues or subtle material differences remain more challenging.

Overall, the thesis demonstrates that PVDF-based tactile sensorization can support both data-driven inference of object properties and feedback-oriented, real-time tactile information delivery on upper-limb prosthetic-hand platforms. By combining probabilistically interpretable hardness inference with scalable distributed sensing and multichannel electrotactile feedback, this work provides a foundation for future sensory restoration strategies and tactile-informed prosthesis control.

Keywords: tactile sensing; electronic skin; object hardness discrimination; machine learning; deep learning; real-time processing; electrotactile feedback; upper-limb prosthetics.

Contents

List of Figures

List of Tables

Acronyms

CNN	convolutional neural network
HR	high resolution
OA	overall accuracy
OSM	OpenStreetMap
PCA	principal component analysis
FD	Finger Detection
SiD	Size Detection
PD	Closed hand position Detection
SpD	Speed Detection
HD	Hardness Detection

CHAPTER 1

Introduction

1.1 Motivation

The human sense of touch is a fundamental channel for interacting with the physical world. Mechanosensory neurons embedded in the skin encode contact events and convey information such as pressure, vibration, and pain to the central nervous system, enabling both perception and protective reflexes [4, 5]. During everyday manual actions, tactile signals from the fingertips are tightly coupled with motor control: they inform grip-force regulation, manipulation stability, and rapid corrective responses [6]. Beyond functional control, touch also contributes to a coherent bodily self-representation, supporting the feeling that the hand is part of one's body [7, 8]. This close coupling between sensation and action motivates the broader goal of equipping artificial hands with the ability to sense contact and to transform that information into actionable or perceivable cues.

In this context, upper-limb prosthetic hands aim to restore autonomy in activities of daily living, yet the restoration of dexterous motor function alone is not sufficient. Most users still rely heavily on vision to compensate for missing somatosensory cues, which increases cognitive burden and can limit manipulation performance in realistic environments [9, 10]. User studies and surveys consistently report that sensory feedback is a priority requirement for upper-extremity prostheses and an important factor shaping long-term acceptance [11, 12]. Accordingly, sensory-enabled prostheses are increasingly viewed as a necessary step toward more functional and embodied artificial hands [8, 9]. From a systems perspective, enabling "prosthetic touch" requires at least three coupled components: (i) tactile sensing on the prosthesis, (ii) processing/decoding of tactile data into meaningful variables (e.g., contact location, force, or object properties), and (iii) a feedback interface that delivers usable information back to the user [13, 14, 9]. Feedback interfaces include invasive solutions (e.g., peripheral nerve stimulation) that can evoke localized sensations referred to the missing hand [15, 8], and non-invasive sensory substitution approaches (e.g., vibrotactile, mechanotactile, or electrotactile stimulation) that are attractive for safety and ease of adoption [10, 16, 14]. Importantly, even when feedback is available, the usefulness of the information depends on how reliably tactile data can be captured and interpreted under realistic interaction variability.

A core technical barrier in sensory prostheses is therefore robust *sensorization* of the hand surface. To approach the spatial richness of biological touch, distributed tactile sensing arrays ("electronic skin") are required, ideally combining mechanical compliance, conformability, durability, and sufficiently dense spatial sampling [17, 18, 19]. Over the last two decades, advances in materials and fabrication have significantly expanded what electronic skin can provide for interactive robots and wearable platforms [20, 19], and flexible sensor technologies have been extensively reviewed and categorized with an emphasis on their suitability for conformal, lightweight, and wearable deployment [21]. At the same time, tactile sensors based on diverse transduction principles—including capacitive, piezoresistive, piezoelectric, optical, and magnetic approaches—have been proposed for robotic and biomedical applications [22, 23, 24]. However, many technologies remain challenging to deploy on prosthetic hands because practical systems must satisfy constraints on thickness, wiring complexity, robustness to environmental factors, and integration with existing hand mechanics [18, 25]. In other words, the challenge is not only achieving sensitivity in the lab, but achieving stable sensing when the sensor becomes part of a real prosthetic embodiment that must withstand repeated grasping, manipulation, and day-to-day handling.

In parallel, the readout chain (signal conditioning, acquisition, and digitization) strongly influences what information can be extracted from tactile sensors, particularly for piezoelectric polymers. Interface electronics for PVDF-based tactile sensors have therefore been actively studied, with attention to noise, bandwidth, and practical integration [26, 27]. This electronic layer is not merely a support component: it determines whether tactile sensing is usable for downstream decoding in real time, and whether the overall system can be scaled beyond proof-of-concept demonstrations. Consequently, tactile sensorization must be considered jointly with the electronics and the data pipeline, rather than as independent blocks.

Even with high-quality tactile sensing hardware, translating raw signals into actionable information remains non-trivial. Tactile data are often high-dimensional, noisy, and strongly dependent on contact conditions, loading profiles, and interaction dynamics [18, 28]. For that reason, tactile data decoding spans a spectrum from classical signal processing (e.g., denoising, drift removal, event detection, slip indicators) [29] to learning-based inference of object properties and interaction states [30, 28]. Among object properties, *hardness* (or stiffness-related discriminability under interaction) is particularly relevant for functional grasping and safe manipulation. Humans perceive hardness using interaction-dependent cues (e.g., transient forces and vibration during tapping or indentation), and perception can vary across moderate stiffness levels that produce similar responses [31, 6]. In robotics and tactile AI, hardness estimation and classification have been explored using both engineered features and deep learning approaches, including shape-independent strategies under controlled sensing conditions [32, 30]. However, transferring such ideas into sensory prosthetic hands requires addressing additional constraints: *real-time operation*, robustness to speed/load variability, limited on-device resources, and the need for interpretable outputs that can support feedback delivery rather than offline analysis only [10, 9]. These constraints become particularly important when the goal is not just classification accuracy in isolation, but the extraction of tactile information that can be acted upon quickly and reliably during manipulation.

Despite major advances, several gaps remain before tactile sensing and feedback can become practical and beneficial for everyday upper-limb prosthesis use.

1: Practical distributed sensorization of prosthetic hands. Distributed tactile sensing must remain thin, compliant, and robust while maintaining stable electromechanical coupling and scalable interconnects. Although e-skin concepts are mature in robotics, prosthetic integration still requires careful design choices to balance spatial coverage, reliability, maintainability, and wearability [18, 19, 25].

2: Decoding pipelines that work *online* under realistic conditions. Many tactile classification studies report high accuracy under controlled protocols, yet real prosthetic interactions involve variability in contact timing, indentation speed, loading, and incidental disturbances. Robust online pipelines therefore require signal processing for event detection and drift/noise handling, plus learning models that can generalize across conditions while meeting latency constraints [30, 29, 10].

3: Closing the loop with feedback. For feedback to be useful, the system must output variables that can be mapped to stimulation in a stable and interpretable way. This includes not only point estimates (e.g., predicted hardness class) but also confidence/uncertainty measures that can support confidence-aware strategies [14, 8]. Invasive interfaces can provide high-fidelity sensations but face surgical and translational barriers; non-invasive interfaces are more deployable but require careful coding and stimulus design to remain informative and comfortable over time [16, 10, 9].

Motivated by these gaps, this thesis targets the development and validation of a tactile sensing framework for upper-limb prosthetic hands that progresses from *localized biomimetic sensing* to *hand-level sensorization*, and finally to *distributed sensing with feedback-ready outputs*. The experimental platform is centered on prosthetic-hand embodiments enabling realistic manipulation contexts [33], and on an e-skin and electronics stack designed for piezoelectric polymer sensing, where interface design and noise-aware processing are integral [26, 27].

A key objective is to demonstrate that tactile signals acquired from piezoelectric sensor arrays can be decoded to infer meaningful object-related information (with emphasis on hardness discrimination) under interaction variability, using machine learning and deep learning approaches designed for online operation [30, 32].

In this thesis, the motivation is to move beyond that limitation by demonstrating methodologies on a real prosthetic-hand platform, where the sensing system can be mounted and the hand can physically interact with objects in human-like scenarios [33]. This distinction is central because it directly affects integration constraints (geometry, cabling, robustness) and the realism of the tactile data distribution that the decoding and feedback pipelines must handle.

Beyond decoding, the thesis aims to demonstrate how distributed tactile information can be transformed into *feedback-ready* signals that support end-to-end sensory restoration through non-invasive stimulation interfaces, leveraging established principles for wearable haptics and tactile displays [14, 13, 16].

In this way, the overall motivation is not only to improve tactile sensing hardware, but to validate a complete chain that connects sensorization, real-time decoding (using ML/DL and signal processing), and feedback-oriented outputs in prosthetic-hand conditions that are as realistic and transferable as possible.

1.2 Contributions

This thesis advances upper-limb prosthetic sensing and sensory feedback by developing and validating a tactile sensing-to-feedback pipeline that progressively increases realism and system scope: from a biomimetic fingertip sensor used in controlled indentation experiments, to prosthetic-hand hardness inference under less controlled grasps, and finally to a full distributed tactile sensing and electrotactile feedback interface evaluated in human-subject experiments. The main contributions are grouped into three complementary themes: (i) tactile sensing and sensorization, (ii) tactile data processing and hardness inference, and (iii) end-to-end sensory feedback for prosthetic hands.

1.2.1 Progressive tactile sensing and sensorization for a prosthetic hand

A central contribution of this thesis is the step-by-step development and validation of a tactile sensing system that evolves from a biomimetic fingertip prototype toward full-hand sensorization on the Hannes prosthetic hand.

> **Chapter 3 (biomimetic fingertip sensing, controlled validation):** Integration and experimental validation of a PVDF/P(VDF-TrFE)-based biomimetic finger-cap sensor (8 sensing points) together with a dedicated readout circuit, and systematic characterization through a controlled Cartesian-robot indentation protocol that varies indentation speed and target load to capture realistic contact dynamics.
>
> **Chapter 4 (prosthetic-hand sensorization for hardness inference):** Transfer of the sensing concept from a controlled laboratory indentation setup to a less controlled prosthetic-hand context, using the Hannes hand as the experimental platform to evaluate hardness discrimination from tactile sensing during grasping, and to quantify the robustness of tactile inference across grasp conditions.
>
> **Chapter 5 (distributed sensorization on the prosthetic hand):** Realization of a distributed tactile sensing system on the Hannes hand using a multi-sensor e-skin architecture that extends sensing coverage beyond a single sensing site, enabling spatially structured interaction sensing (fingertips/palm) and supporting subsequent closed-loop mapping to distributed stimulation.

1.2.2 Tactile data processing and hardness inference under increasing variability

A second contribution is the design and evaluation of tactile processing pipelines that demonstrate how hardness information emerges over time, how model choice affects robustness, and how probabilistic outputs can be exploited for reliable online decision-making.

> **Chapter 3 (feature-based ML + incremental inference analysis):** Development of two hardness-discrimination pipelines: (i) a reduced single-sensor approach using PCA-ranked statistical features and classical classifiers, and (ii) a multi-sensor incremental framework that explicitly studies hardness discrimination as a function of the observed portion of the signal using cumulative windowing strategies, complemented by probabilistic analyses (Top-k, calibration, confidence structure) to support confidence-aware decision rules.
>
> **Chapter 4 (CNN vs. feature-based baselines + scenario-wise robustness + online criteria):** Benchmarking of classical ML models against an end-to-end CNN trained on raw windows across multiple window durations, followed by scenario-wise evaluation across grasping conditions (target contact pressure and grasping frequency). In addition, the chapter proposes and validates confidence-gated online decision rules (event-level voting and detection latency under a confidence threshold), enabling realistic performance assessment in real-time operation.

1.2.3 End-to-end sensory feedback system for upper-limb prosthetics

A third contribution is the end-to-end distributed sensing to interpretable, spatially organized tactile feedback on the user, validated through psychophysical experiments.

> **Chapter 5 (distributed sensing → distributed electrotactile feedback):** Design and validation of an end-to-end sensory feedback system that maps tactile events acquired from a distributed e-skin on the Hannes prosthetic hand to spatially structured electrotactile stimulation delivered through flexible matrix electrodes. The full pipeline is evaluated in human-subject experiments across multiple perceptual

tasks, demonstrating the feasibility of high-channel, real-time, non-invasive tactile feedback for upper-limb prostheses.

1.3 Thesis outline

The remainder of this thesis is organized as follows.

Chapter 2 reviews the state of the art in upper-limb prosthetic sensing and sensory feedback, covering tactile sensing modalities, integration constraints, tactile signal processing and inference, and non-invasive feedback interfaces. The chapter highlights open challenges in robustness, real-time operation, and the translation from laboratory demonstrations to prosthetic-ready systems.

Chapter 3 presents the development and experimental validation of a biomimetic PVDF-based fingertip sensing system and its use for hardness discrimination in a controlled Cartesian-robot indentation setup. The chapter introduces two activities: (i) feature-based hardness classification using a reduced single-sensor pipeline, and (ii) an incremental multi-sensor framework that analyzes how classification performance evolves as a function of the observed portion of the tactile interaction, complemented by probabilistic confidence analysis.

Chapter 4 investigates hardness classification in a less controlled setting using a prosthetic hand. Classical machine learning models and a CNN are evaluated across window durations, followed by scenario-wise analysis across grasp conditions. The selected CNN is deployed in a real-time pipeline (PC-based) and evaluated online using event-level decisions and confidence-gated detection with latency analysis.

Chapter 5 presents an end-to-end sensory feedback system for upper-limb prosthetics, integrating distributed tactile sensing on the prosthetic hand with distributed electrotactile stimulation on the user. The complete pipeline is evaluated in human-subject experiments across multiple perceptual tasks designed to probe complementary aspects of tactile information relevant to manipulation.

Chapter 6 concludes the thesis by summarizing the main findings, discussing limitations, and outlining directions for future work toward robust, scalable, and clinically relevant tactile sensing and sensory feedback for prosthetic hands.

CHAPTER 2

State of the Art

2.1 Introduction

Tactile sensing and sensory feedback are central to achieving stable grasping, safe physical interaction, and a feeling of embodiment in robotic and prosthetic hands. Despite major advances in mechatronic design and control, many systems still operate with limited information about *when*, *where*, and *how* contact occurs, which constrains manipulation performance and increases the user's reliance on vision. In the context of upper-limb prostheses, closing the loop with somatosensory feedback is widely recognized as a key step toward more natural and effective use [34, 35]. More broadly, drawing inspiration from the structure and bandwidth of human cutaneous afferents [36, 37], artificial tactile systems combine transduction materials, electronics, and signal processing to transform physical interactions into informative features and actionable feedback. Recent progress in printed and flexible sensor technologies has further enabled conformable, scalable sensorization over complex hand geometries [38, 39], while embedded processing and closed-loop feedback interfaces have moved tactile sensing from purely offline analysis toward real-time operation [40, 41, 42].

However, building a practical tactile-enabled prosthetic/robotic hand remains a system-level challenge: the sensing modality and packaging must be robust and mechanically compliant; the data pipeline must handle noisy, non-stationary signals under variable contact dynamics; and the extracted tactile information must be encoded through a feedback interface with sufficient bandwidth and repeatability to be learnable and useful in functional tasks. These considerations motivate a structured view of the literature that spans (i) tactile sensor technologies, (ii) tactile data processing (signal processing and machine learning), (iii) representative tactile classification and control applications, and (iv) feedback techniques used to convey tactile information to the user.

Accordingly, this chapter reviews tactile sensing systems with an emphasis on prosthesis- and teleoperation-oriented solutions. It first overviews tactile sensor technologies, focusing on piezoelectric, resistive, capacitive, and optical approaches (Section 2.3). It then summarizes tactile data processing pipelines, covering signal processing and machine-learning-based interpretation (Section 2.4), including representative learning strategies used for tactile time-series [43, 42]. Next, it surveys tactile-sensing applications relevant to this thesis hardness classification, as well as grasp control and slip detection (Section 2.5). Finally, it reviews major feedback modalities (vibrotactile, electrotactile, mechanotactile, and kinesthetic) and discusses how their constraints influence encoding strategies for delivering tactile information (Section 2.6). The chapter concludes by highlighting the key gaps that motivate the thesis approach: scalable PVDF-based sensorization, real-time processing, and distributed non-invasive feedback suitable for end-to-end closed-loop operation.

2.2 Tactile Sensing Systems

Sensors for robotic hands transduce various modalities of tactile stimuli, aiming at recreating naturalistic perception. Today's tactile sensing systems encounter many challenges that limit their integration in robotic hand systems (e.g., prosthetic systems), such as design issues, spatial distribution, low signal-to-noise ratio (SNR), crosstalk, wireless communication, and the lack of signal processing and machine learning methods to encode the acquired data from tactile sensors [44]. The main design requirements are summarized in table 2.1.

It is expected that artificial tactile sensors demonstrate small spatial resolution (≥ 1mm for fingertips, 5 mm for hand palm), high sensitivity varying from 0.01 to 10 N, which extends along with the tactile frequency range (< 1Hz - 1kHz), low hysteresis, fast and linear response (less than 1 ms), wide dynamic range, and high reliability. Furthermore, it needs to exhibit high electromechanical bandwidth to detect fast events (e.g., incipient slip) and customizable shapes of e-skin patches and sensor numbers. Additionally, low cost, low power consumption, and scalability are major factors for the robotic application.

Table 2.1: Design requirements for tactile sensing system

Design criteria	**Character guideline**
Detectable force range (Dynamic range)	0.01 N–10 N
Spatial resolution	≥ 1 mm for small sensing areas (e.g., fingertips), ≥ 5 mm for large sensing arrays (e.g., limbs, torso, etc.)
Sensor frequency bandwidth (sensor response time)	0.1 Hz–1 kHz, about 1 kHz (1 ms)
Mechanical sensing detection capability	Normal and shear forces; vibrations
Sensor System characteristics	**Mechanical:** Flexible, stretchable, conformable, and soft, robust, and durable. **Electrical:** Low power, minimal wiring and cross-talk, and electrically and magnetically minimal sensitivity.
Sensor response	Monotonic, not necessarily linear, low hysteresis, stable and repeatable
Temporal variation	Both dynamic and static

2.3 Tactile Sensors

Enabling tactile sensation particularly in upper limb prostheses is still not as mature as that in other fields (such as robotics, touch screens, etc.), yet the achievements could be adopted into this field. Given that grasping is one of the major functions of hands, most studies of tactile sensing focus on grasp force or pressure to prevent slip and achieve a stable grasp. The measured characteristics of touch, however, can be not only force and pressure, but also stiffness, texture, or shape. Thus, different sensing techniques are desired to be synthesized to realize a human-like tactile sensing system [45].

The following is a review of the available tactile sensing techniques which have the potential to be applied in robotic hands and specifically in prosthetic hands, namely, resistive sensors (such as strain gauges and piezoresistive), capacitive sensors, piezoelectric sensors, and optical sensors. Table 2.2 summarizes the advantages and disadvantages of each of the aforementioned sensing techniques.

Table 2.2: Characteristics of different tactile sensing techniques used for e-skin fabrication

Tactile sensors	Working principle	Advantages	Disadvantages
Piezoresistive	Its resistance varies with the deformation caused by applied force.	Simple electronics; High sensitivity; Resistant to interference	Hysteresis; Temperature Sensitivity; Fragile and Rigid; High power consumption
Capacitive	Its capacitance varies with the deformation caused by applied force.	Sensitivity of small force change; Reliability; Dynamic and static force measurement; Low temperature sensitivity; Low power consumption.	Limited spatial resolution; Noise sensitivity; Complex electronics; Cross-talk between elements; Hysteresis
Piezoelectric	An electric voltage will be produced when a force is applied to it.	No need for power supply; High reliability; Fast dynamic response; High sensitivity; High accuracy	Low spatial resolution; High temperature sensitivity; Inability to sense static value
Optical	The intensity or the spectrum of light varies with the applied force.	Immune to electromagnetic fields; High spatial resolution; Wide sensing range; Good reliability	Fragile and rigid; Large size

2.3.1 Piezoelectric sensors

Piezoelectric materials transfer mechanical stress into an electrical potential. Piezoelectric sensing is one of the few sensing techniques that do not require a power supply, which is considered an outstanding advantage. Besides, it also exhibits high sensitivity, reliability, and fast dynamic response. It's wide response range of 0 to 1 kHz makes it a good choice for vibrations measurement [46].

Various piezoelectric materials can be used for constructing piezoelectric tactile sensors. The most widely used one is polyvinylidene fluoride (PVDF). PVDF has many advantages, like mechanical flexibility, dimensional stability, high piezoelectric coefficients, formability into very thin sheets (5 μm), and a relatively low price. Moreover, it has a fast and accurate response to high-frequency vibrations. Polymers such as PDMS have been used in mesa micro-structures to achieve flexible and sensitive sensors [47]. Another promising piezoelectric material is zinc oxide (ZnO) nanotransducer because of its high flexibility and biocompatibility [48][49]. ZnO is proposed to be a good candidate material for pressure and temperature sensor to be applied to prosthetic limbs.

During the past years, piezoelectric sensors have been used for the detection of slip [48], texture [50][51], roughness [52] and stiffness [53]. Thus, many surface characteristics can be completely determined using piezoelectric sensors. However, piezoelectric tactile sensors also have inherent drawbacks, the greatest one of which is the inability to measure static contact forces. Moreover, robustness and sensitivity to temperature are lacking.

2.3.2 Resistive sensors

Resistive sensors measure internal changes in resistance due to an externally applied force. Resistive sensors have been used to detect the normal force, shear force, lateral strain, and bending strain through the conversion of such mechanical changes into an electrical resistance [54][55].

Certainly, sensors that are more useful for robotic hands are those that are more sensitive to a range of low-impact forces that would typically be experienced at the fingertips. Jorgovanovic et al. presented the static and dynamic characterization of piezoresistive sensors used for detecting the positions of prosthetic finger joints [56]. Kane et al. [57] proposed a piezoresistive stress sensor array with high spatial resolution comparable to the human dermis. It exhibited a high potential for dexterous manipulation applications. Various applications with piezoresistive tactile sensors can also be found in stress and force measurement [57], stiffness of soft tissues detection [58], and for fingertip sensing [59] etc.

2.3.3 Capacitive sensors

Generally, a capacitive sensor consists of a dielectric material sandwiched between two parallel conductive layers. The capacitance between the two layers varies with the deformation of the dielectric material as a result of the force applied to the sensor. The capacitive sensor is considered the most sensitive sensor for detecting small force changes. Moreover, it exhibits high sensitivity, robust performance, a large dynamic range, temperature sensitivity, and low power consumption [60]. It can be used for both dynamic and static force measurement. Additionally, their sensitivity to noise leads to relatively complex electronics for noise filtration. Several designs for pressure detection, shear force sensing, and texture recognition have been proposed in the literature [22]. The capacitive touch sensor array was integrated into a prosthetic hand thumb finger in [61]. Capacitive sensors have also been deployed for multi-axis force measurement in gripping and object manipulation, texture recognition [60], shear sensing [23], and touch screen applications [62], among others. Another capacitive tactile sensor was presented for gripping force measurement with a sensor range of 0-3000 mN [61].

2.3.4 Optical sensors

Optical sensors have been used for tactile signal transduction, measuring optical variations across semitransparent media due to physical deformation upon contact and pressure [24]. Optical sensors maintain high sensitivity and resolution while being immune to electromagnetic interference from nearby sources [23]. These sensors are subject to less hysteresis and time response than other types of devices due to the immediate response of light intensity to strain in the device. Recently, a prosthetic finger was developed using a sensor with a looped optical waveguide [63].

2.4 Tactile Data Processing

2.4.1 Signal Processing

The quality of the data used to train machine-learning models is critical, since noise, artifacts, and irrelevant variability can directly degrade classification performance. Consequently, tactile signals are commonly pre-processed to reduce noise and to extract informative representations and features. In addition to improving robustness, feature extraction can reduce the computational and memory resources required for downstream learning and inference.

A widely used approach is to transform the raw time-domain signal into a frequency-domain representation using the Fourier transform (FT). In practice, the discrete Fourier transform is typically computed via the fast Fourier transform (FFT), which enables efficient processing of temporal sensor signals. Frequency-domain signatures can capture

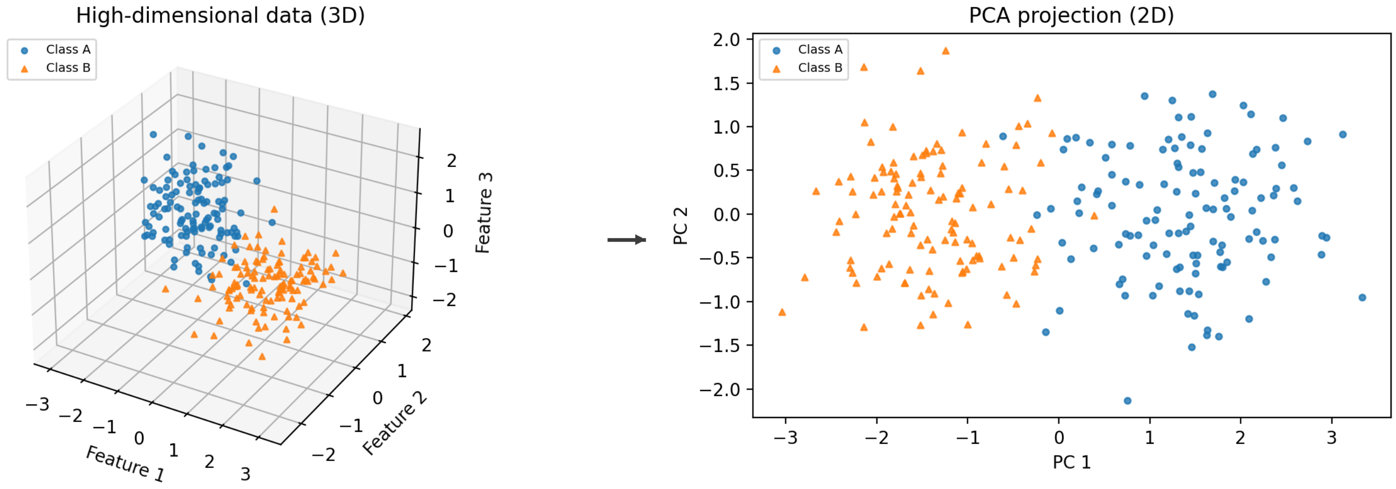

Figure 2.1: PCA example illustrating dimensionality reduction from 3D to 2D (generated using the data and code in [1]).

interactions such as micro-vibrations induced by contact and sliding, and have been used to characterize surface features by relating spectral changes across tactile sensor signals [64, 65]. These spectral descriptors can be used directly as features for machine-learning algorithms [66, 67]. Kursun and Patooghy [68] investigated FFT-based feature extraction using a window length of 256 samples and evaluated multiple learning algorithms; however, they reported that real-time embedded implementation was hindered by code-size and memory constraints. More generally, comparable classification performance can often be achieved using time-domain and frequency-domain representations, depending on the task and model capacity [69, 70, 71].

When FFT-based descriptors are not sufficiently expressive, time–frequency methods such as the discrete wavelet transform (DWT) are frequently employed [72, 73]. DWT decomposes the signal into multiscale components by passing it through paired low-pass and high-pass filter banks, yielding sub-band representations that can be used to form feature vectors [73]. In the context of surface roughness discrimination, Qin *et al.* [67] applied DWT to PVDF tactile signals and selected discriminative wavelet-derived features for classification. For texture recognition using artificial fingers, Hanif *et al.* [74] extracted FFT-based tactile features during exploratory movements; a key limitation of this class of approaches is that texture decoding can depend on controlling and/or

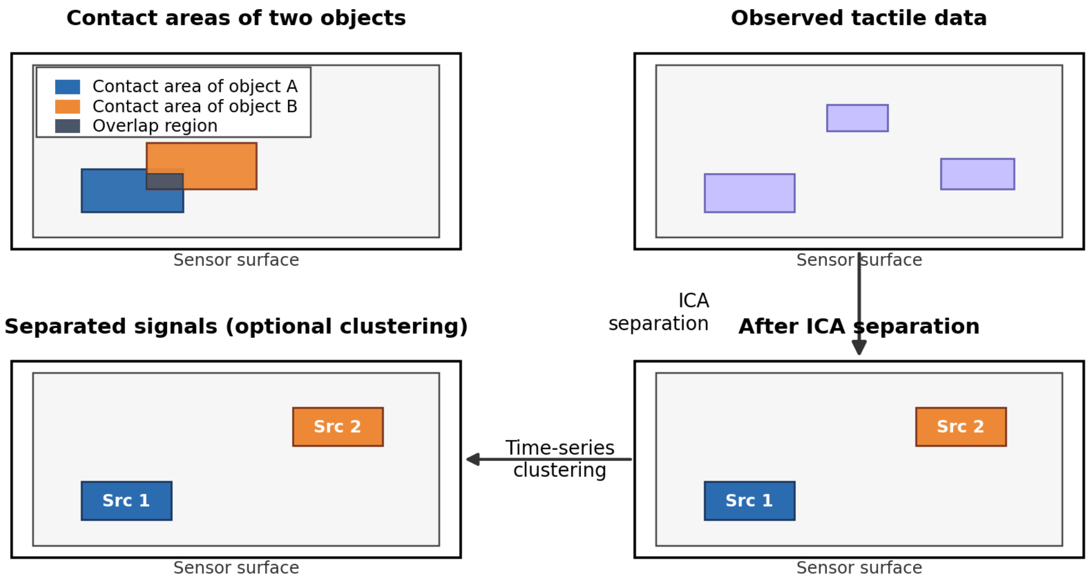

Figure 2.2: Conceptual procedure for separating tactile information from multiple sources using ICA, optionally followed by time-series clustering (concept adapted from [2]).

knowing the sliding conditions (e.g., velocity and contact trajectory), which may not be available in unconstrained interaction.

In parallel to spectral and wavelet descriptors, classical time-domain statistical features are widely used due to their low computational cost. Typical examples include mean, standard deviation, variance, skewness, kurtosis, extrema, and signal power, among others [75, 76]. Pre-processing may also include reducing redundant samples via sub-sampling strategies prior to feature computation [77], which can further decrease processing load for embedded tactile pipelines.

Beyond hand-crafted features, dimensionality reduction and feature extraction techniques are often applied to compact high-dimensional tactile data. Principal component analysis (PCA) is a foundational tool in multivariate data analysis [78]. PCA projects data into a lower-dimensional subspace that preserves maximal variance, enabling a representation of data in $\mathbb{R}^{n-k}$ from an original $\mathbb{R}^n$ (for positive integers n and k) with minimal information loss. Figure 2.1 illustrates an example of reducing a 3D dataset into a 2D representation (the figure was generated using the data and code in [1]). PCA has been applied in tactile and material recognition pipelines, for example to re-

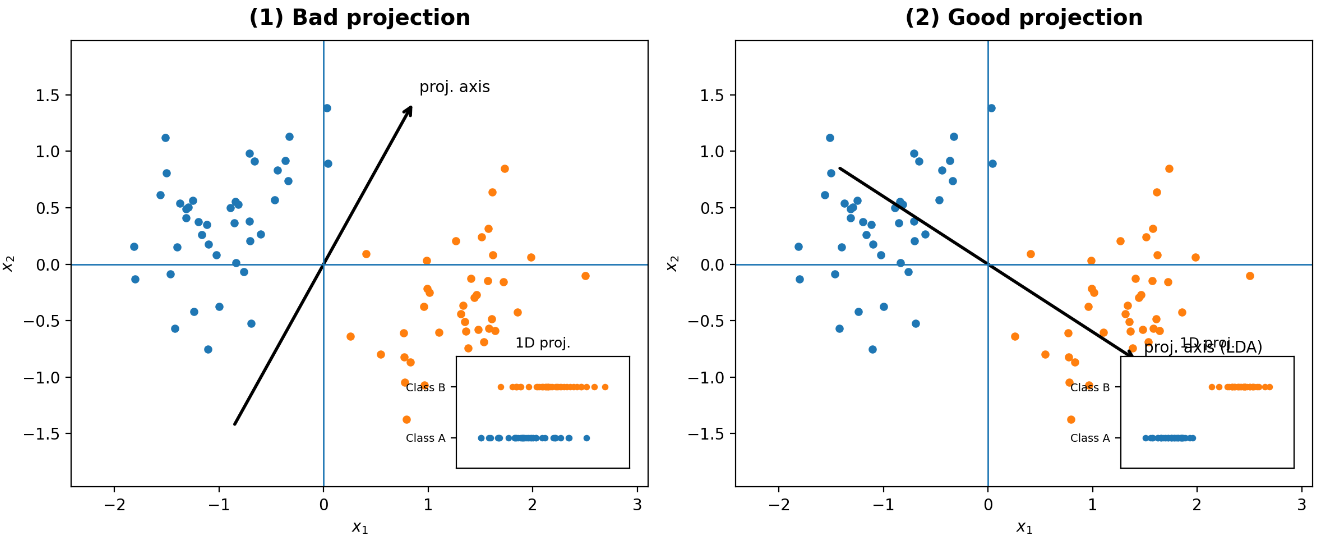

Figure 2.3: LDA intuition: a poor projection direction can yield overlapping class projections, whereas a discriminative direction increases separability in the projected space [3].

duce dimensionality within individual sensing modalities of the BioTac sensor prior to training a neural network [79], and to identify relevant features from thermal-property datasets for material classification [80]. PCA-based compression has also been used in tactile arrays to retain the most informative pressure patterns [75], and in PVDF-based fabric evaluation where FFT-derived descriptors are compressed before classification [81]. Nonlinear extensions include kernel PCA [82], which enables nonlinear feature extraction through kernelized eigen-decomposition; it has been applied, for instance, to low-resolution tactile image recognition in robotic assembly [83]. Other nonlinear reductions include local PCA [84], which has been combined with learning methods for tactile-based object classification [85].

Independent component analysis (ICA) can be viewed as a related linear technique that seeks statistically independent components rather than orthogonal directions of maximal variance [86]. ICA is often introduced through the "cocktail party" analogy, where the objective is to separate multiple mixed sources into independent signals [87]. In tactile perception, spatial ICA has been used to separate tactile information arising from multiple simultaneous sources, for example in combination with time-series clustering [2]; Figure 2.2 provides a conceptual overview of this procedure.

Finally, linear discriminant analysis (LDA) is another classical method for feature extraction and dimensionality reduction that explicitly optimizes class separability. LDA finds a linear combination of features that maximizes between-class separation relative to within-class scatter [3]. Unlike PCA, which is unsupervised, LDA uses class labels and is therefore directly aligned with classification objectives. Figure 2.3 illustrates how different projection directions can lead to poor or strong separability in the projected space. LDA has been used, for example, to classify deformable versus non-deformable surfaces from tactile images [88], and has also been explored for surface texture discrimination [89]. In a broader object-recognition setting with the BioTac sensor mounted on a robot, Hoelscher *et al.* [90] compared several tactile feature representations (including PCA-based variants and physically motivated features) and demonstrated reliable interactive recognition across a large set of objects.

2.4.2 Machine Learning

Machine learning (ML) is a research field focused on developing methods that *learn* from data that is, methods that improve their performance on a task through experience rather than being explicitly programmed. It is commonly viewed as a sub-area of artificial intelligence. In ML, models are trained on example observations (training data) to produce predictions or decisions on unseen data [70].

When building an ML model, one must choose among many possible algorithms; however, strong performance in one domain does not guarantee strong performance in another. The well-known *no-free-lunch* theorem formalizes this idea: even if an algorithm performs well for a particular class of problems, its performance cannot be assumed to transfer to a new problem without empirical evaluation [91].

In tactile perception and object-property recognition, both neural-network approaches ranging from shallow networks to deep architectures such as ANN, CNN, RNN, and transformers and classical ML methods such as SVM, KNN, extreme learning machines (ELM), Naive Bayes, and decision trees have been widely used to classify and identify properties of touched objects [92]. Finally, across prosthetics-related sensing and perception studies, supervised learning with labeled data remains the dominant paradigm [93].

Table 2.3: Machine and Deep Learning Algorithms for Tactile Data Applications

Application	Learning Algorithm	Examples
Hardness/softness detection	Support vector machine, K-nearest neighbors, Decision Tree, Naive Bayes, Convolutional neural network, recurrent neural networks, K-medoids	[94],[95],[75], [96], [97],[98]
texture/roughness detection	Artificial neural network, Support vector machine, Naive Bayes, K-nearest neighbors	[99],[100],[101], [102],[103],[104], [105]
Slip and Grasp Detection	Neural networks, Support vector machines, K-nearest neighbors, logical regression, Convolutional neural network, K-Means	[106],[107],[108], [109],[110],[111]

2.5 Tactile sensing classification Applications

Machine learning algorithms are designed to emulate human intelligence by learning from the surrounding environment [112]. ML algorithms can extract complex, nonlinear input-output relationships. They are trained using a set of examples, where each example is described by a group of informative features. ML algorithms can support intelligent and predictive systems that can make accurate decisions on unseen data. In this perspective, several works in the literature witness the adoption of machine learning algorithms for classification/regression problems in robotic tactile sensing systems, including: object recognition, object properties classification (e.g., Hardness, shape, and texture, etc.), contact properties control (e.g. stability control, grasp control, and manipulation), and contact events detection (e.g. detection of slippage), using normal and shear force sensing, in addition to vibration detection. Moreover, for tactile information retrieval there is a need for physical contact to be made with the object, material or human. Tactile sensors supply various attributes about objects they contact that vision sensors simply cannot, such as hardness, texture, weight, etc. Table2.3 reports the commonly used ML algorithms with respect to the type of extracted tactile data.

2.5.1 Hardness Classification

A common characteristic of an object that can be represented by tactile images is hardness. The work in [113] described the hardness of the object as the ratio between the force applied to an object and the resulting displacement of the object. The authors in [75] integrated a tactile sensing system into a two–finger robotic gripper for the hardness classification of fruits and vegetables. The compression test technique was adopted to measure and label the hardness of the objects. Tactile sensors were mounted on both clamps of the gripper, and data was collected by grasping each object multiple times. Two models, i.e. support vector machine (SVM) and k–nearest neighbor (KNN), were trained to solve a four–class hardness classification problem. As a result, the models achieved a high classification accuracy (94.37%). In [94], piezoresistive tactile sensors were integrated into a robotic gripper to explore the hardness of seven different types of fruit by squeezing them. The tactile information was presented via a temporal sequence of images that encode the pressure applied to the taxels. Each tactile image is an array of 64 values (8 x 8). The mean and standard deviation were extracted as features for each tactile image (one frame) and then used to train a kNN classifier to classify the fruits. Using k=1, the model achieved a classification accuracy of 92.86% [94]. For vegetable sorting, a two-finger robotic gripper was equipped with a tactile sensing system to discriminate between green, moderate, and ripe tomatoes using decision tree (DT) and Naive Bayes (NB) algorithms [114]. Tactile data was collected while grasping objects. The ML models were deployed into a PIC32 microcontroller to achieve real-time softness classification. An accuracy of 90% and 85% was achieved using DT and NB, respectively. In [32], the authors proposed a CNN and a recurrent neural network (RNN) to estimate the hardness of objects with different shapes using a GelSight tactile sensor [115]. The sensor was integrated into the fingertip of a robot and data was recorded as the object was grasped. The network was able to predict well the hardness of silicone samples with similar shapes in the data set, regardless of the loading conditions; however, for objects with rigid surfaces, the model was not able to estimate their hardness wel. In [97], a skin sensor covering the forearm of a humanoid robot ’Cody’ was used for a classification problem that involves classifying 18 objects based on hardness (soft or rigid) and state (fixed or movable) using a KNN classifier. For k=2, a classification accuracy of 80% was achieved. An FPGA-based tactile system was mounted on a

Cartesian robot where tactile data was collected as objects were squeezed [98]. Two algorithms were considered for the hardness classification of nine objects: K-medoids and KNN. Each of the two algorithms achieved a classification accuracy of up to 86.7%. For object classification based on shape, the authors in [116] integrated a piezoresistive Tekscan tactile sensing system into a robotic finger. Pressure map images were collected upon grasping four objects of different shapes. The authors developed a novel algorithm to extract features from the pressure maps, which were used to train an artificial neural network (ANN) to classify the shape of objects. As a result, the model achieved a 90% success rate.

2.5.2 Texture Classification

According to [94] , roughness is the most studied material property in the context of tactual perception. Authors in [99] investigate the ability of a biomimetic fingertip to classify between different fabrics based on the texture of their surface alone using ANNs. The fingertip contains five accelerometers and eight single point force sensors. For data collection, the fingertip was moved in an exploratory motion to collect data from eight force sensors and five accelerometers. The experiment was carried out on seven types of materials made of fabrics. Using ANN, the model achieved 85% classification accuracy based on texture. However, some textures were frequently misclassified such as polyester and acrylic. In [100], authors proposed a tactile sensing system based on four piezoelectric sensors to detect the surface roughness of fruits and vegetables using a support vector machine (SVM) algorithm with a radial basis function kernel. In addition, CNNs were employed for feature extraction from spatially distributed tactile sensors. The model achieved a 91% success rate solving 4-class classification problem. In [101] multiple ML algorithms were implemented to classify three different materials (denim, a photo and tape) using texture related tactile information from their developed sensors. Two ML algorithms, a multi-layer ANN and NB, were fed with data input extracted using discrete wavelet transform. It was found that the ANN performed the best by achieving 91% classification accuracy between objects of different materials, while the NB achieved 67% accuracy. However, neither ML methods was able to distinguish accurately between the two similar materials (i.e., photo and tape). Using SVM, an autonomous humanoid robot equipped with artificial skin on both arms obtained a recognition rate up to 100% based on texture, with 70% objects categorization ability in a setup that involved ten different objects[102] . In [103], the PVDF piezoelectric sensor was integrated to a robotic fingertip for object textures recognition. Tactile data generated from the sensor were used to feed different ML algorithms including SVM, KNN, and artificial neural network (ANN). The SVM and KNN models performed equally well giving classification rates higher than 90%, unlike the performance of ANN which was lower. In [104] a naive Bayes classifier was used to distinguish textures sensed by sliding a sensorized robotic finger along materials and collecting vibration data. Different textures induce different intensity of vibrations in the sensor. Textures can be distinguished by the presence of different frequencies in the signal. A total of seven materials were tested including sponge, carpet, wood, two tiles of different roughness and two pieces of vinyl of different roughness. The data from the finger was pre-processed and the Fourier coefficients of the sensor outputs were used to learn a classifier for different textures. The NB classifier was capable of predicting textures between dissimilar surface textures with an accuracy of 78%; however, the classifier was unable to distinguish between two types of tiles, as the texture of the surfaces were similar. In [105], Least Square SVM was adopted to discriminate 20 daily used objects based on their texture. A classification accuracy between 70% and 100% has been recorded when using 10 training samples. Bayesian Exploration and Reinforcement Learning have been used to train and validate a discrimination system in [117]. The system was able to differentiate between 10 objects (brick, copper, wood, etc.) with a 90% success rate.

2.5.3 Grasp Control and Slippage Detection

Slip and grasp detection is an another task that can be supported by tactile data using learning algorithms. In [106], a Phantom Omni arm [118] has been equipped with a tactile sensory array of 84 sensor cells to study the translational and rotational movement of an object. The arm was able to hold and recognize a ball with an accuracy of 91.2% using K-Means clustering algorithm. When measurements of vibrations are used to prevent a slippage, then the response time of a sensor becomes crucial. In [107], a robotic system consisting of Barrett arm was equipped with three BioTac sensors to estimate force, and detect and classify slip events for grasped objects. Different ML algorithms were tested including a single-hidden-layer neural network (NN) and 3-layer NN. The best performance was achieved by detecting and classifying a slip with over 80% success rate and within 30 ms. The humanoid robot ARMAR-IIIB [119] was learned to grasp objects using SVM[108]. The grasp was recorded as successful or not by the ability to lift up the object. 77% of the grasps were considered as stable compared to 23% unstable tries. The authors in [109] proposed an intelligent robot grasping system that can grasp objects stably using tactile data from a flexible piezoresistive tactile sensor array. A data collection system was assembled on the robot end-effector. The grasping contact force was collected while grasping several objects and the stability state of each grasping operation was recorded. Three different prediction models, namely support vector classification (SVC), logical regression algorithm (LR), and KNN, were trained to judge the grasping state. It was found that the SVC algorithm-based prediction model yield the best performance by achieving over 98% judgment accuracy for the overall objects, with limited training data. A CNN-based model was trained in [110] to process the tactile information in order to enable successful in–grasp manipulation with untrained daily objects. As a result, CNN effectively handled the tactile information from uSkin sensors. Furthermore, tactile data from 241 distributed tactile skin sensors were used to train feed–forward and deep neural networks to generate a controlled in–hand manipulation of objects of different sizes and shapes [111].

2.6 Feedback Techniques

Sensory feedback is a fundamental component in upper-limb prosthetics because effective object manipulation does not depend solely on feedforward motor command generation, but also on the continuous perception of contact events, force modulation, object slip, hand posture, and interaction dynamics [120, 121, 122]. In the intact human sensorimotor system, these functions are naturally supported by a rich combination of cutaneous and proprioceptive afferent pathways. By contrast, users of prosthetic devices must often rely predominantly on vision, which increases cognitive burden, slows task execution, and may reduce manipulation confidence and embodiment [120, 10]. Consequently, artificial sensory feedback has emerged as a major research direction for improving prosthesis usability, control, and user acceptance.

From a physiological perspective, haptic perception can be broadly interpreted as the integration of two complementary information channels: cutaneous information, originating from mechanical interactions at the skin surface, and kinesthetic information, associated with limb motion, joint displacement, force, and effort perception [123, 124]. This distinction is particularly relevant in prosthetics, where the natural afferent loop is interrupted and must be partially restored using artificial stimulation delivered to the residual limb or to another intact body region. In this context, the design of a feedback system is not only a matter of stimulus generation, but also of selecting an appropriate modality capable of conveying task-relevant information in a compact, wearable, and cognitively interpretable form [120, 121].

Several feedback modalities have been investigated in the literature, including vibrotactile, electrotactile, mechanotactile, kinesthetic, thermal, auditory, visual augmentation, and invasive neural stimulation approaches [120, 10, 14]. However, the present thesis focuses on four non-invasive haptic techniques: electrotactile, vibrotactile, mechanotactile, and kinesthetic feedback. This choice is motivated by three main considerations. First, these four modalities represent the principal non-invasive strategies most commonly adopted in upper-limb prosthetic research for conveying tactile and proprioceptive information without surgery [120, 10]. Second, unlike thermal or purely auditory cues, they directly target the haptic channel and are therefore more suitable for reproducing information naturally associated with touch and object interaction [14, 123]. Third, these modalities offer a meaningful design space in terms of spatial resolution, temporal bandwidth, wearability, and perceptual richness, making them particularly relevant for comparison in the context of wearable artificial skin and prosthetic feedback systems.

In the following, the discussion is therefore intentionally limited to these four techniques. This targeted selection enables a focused comparison between the main non-invasive haptic modalities that are most relevant to the objectives of this thesis, namely the transmission of distributed tactile information acquired by a high-density piezoelectric sensing system toward the user through an intuitive and wearable feedback channel.

2.6.1 Vibrotactile feedback

Vibrotactile stimulation is provoked by mechanical vibration normal or transverse to the skin surface. Mechanical vibration conveys tactile information by modulating vibration frequency, amplitude, duration, timber, or spatial location [125]. The quality of vibrotactile stimulus perception depends on the frequency of the vibration (10Hz to 500 Hz, which corresponds to the bandwidth of the human tactile sense) [126], the body position, and age of the subject [127]. Vibrotactile devices delivering variable pressure on the skin have been employed [128], for instance, it has been implemented to establish proprioceptive communication between the user and the robotic hand. For example, It is suitable for myoelectric prostheses and EEG- based prostheses (e.g. i-Limb myoelectric prostheses Ottobock [129] MANUS [130], Fluid [131] and Smart hands [132]) since no interference with electrical signals.

2.6.2 Electrotactile feedback

Electrotactile (or electrocutaneous) stimulation elicits tactile sensations within the skin by passing low current to stimulate afferent nerve endings in the PNS [126]. Tactile information is delivered through electrotactile stimulation by modulating the electrical components of the pulse waveform i.e. current amplitude (1-20 mA), pulse waveform (monophasic/biphasic, rectangular/sinusoidal), frequency, pulse width, duration of pulse bursts, electrode properties i.e. size(small/large), conducting material and location [10][133]. The variety in the stimulation parameters leads to delivering high bandwidth of tactile information to the user. Electrotactile stimulation can elicit a range of sensations like tingling, itch, vibration, buzz, touch, pressure, pinch, and sharp or burning pain [134] depending on the stimulation parameter and the location. Since there are no moving mechanical parts, the electrotactile systems consume less power, have low weight, produce less noise, and respond faster compared to other tactile feedback systems. Several electrotactile displays have been developed as sensory aids for hearing [135] and vision [136] and can also be used to create perceptual illusions of surface changes [137].

2.6.3 Mechanotactile feedback

Mechanotactile feedback delivers information through direct mechanical interaction with the skin, for example by normal indentation, lateral skin stretch, local pressure, pin arrays, or other forms of controlled cutaneous deformation [120, 14]. In contrast to electrotactile feedback, mechanotactile stimulation physically reproduces a mechanical event at the skin interface and can therefore generate more naturalistic tactile sensations. This characteristic is especially important in prosthetic applications, where naturalness, intuitiveness, and embodiment are central performance criteria in addition to decoding accuracy or control speed [138, 10].

A major advantage of mechanotactile feedback is its intuitive perceptual quality. Because the user feels an actual pressure, skin stretch, or distributed contact, the resulting sensation often resembles natural touch more closely than electrically evoked percepts [14]. Mechanotactile stimulation has therefore been explored to convey grasp force, contact onset, finger position, and distributed pressure in prosthetic and teleoperation systems [120, 10]. Moreover, skin stretch in particular has been shown to be effective for communicating directional or proprioceptive information, since tangential deformation of the skin can induce a clear perception of motion or force direction [139, 140]. Such properties make mechanotactile interfaces attractive when the goal is to increase realism or improve the interpretability of force-related cues.

Nevertheless, mechanotactile systems also present important engineering limitations. They generally require moving parts, contact mechanisms, transmission elements, or miniaturized actuators capable of producing sufficient and repeatable skin deformation. As a result, they are often bulkier, heavier, and mechanically more complex than electrotactile or vibrotactile solutions [120, 14]. Their response bandwidth may also be constrained by actuator dynamics, and increasing the number of independently addressable stimulation points can become difficult when miniaturization, power consumption, and wearability must all be preserved. In high-density applications, the integration of many mechanotactile channels in a lightweight wearable system remains challenging.

Even with these constraints, mechanotactile feedback remains highly relevant because it offers a strong compromise between non-invasiveness and perceptual naturalness. Recent studies have shown that carefully designed mechanotactile interfaces can improve not only task performance but also embodiment and subjective acceptance of prosthetic devices [138]. Therefore, mechanotactile feedback should not be interpreted merely as an alternative to electrotactile or vibrotactile stimulation, but rather as a complementary strategy that is especially valuable when realistic contact rendering and intuitive interpretation are prioritized.

2.6.4 Kinesthetic feedback

Kinesthetic feedback aims to convey information related to limb position, movement, resistance, interaction force, or effort by mechanically acting on muscles, tendons, joints, or skin regions in a way that induces a proprioceptive-like sensation [123, 14]. In prosthetics and wearable haptics, kinesthetic feedback is typically implemented through grounded or wearable devices that apply force, displacement, or skin-stretch cues capable of eliciting a perception associated with motion, force direction, or limb configuration [139, 140, 144].

This modality is of particular interest because natural hand use depends not only on touch at the skin surface, but also on proprioceptive awareness of finger motion, grasp aperture, and interaction force. When this information is absent, prosthesis users may compensate with increased visual monitoring and reduced confidence in manipulation [120, 122]. Kinesthetic feedback therefore addresses an important dimension of sensory restoration, especially for tasks involving graded force control, motion guidance, or perception of hand state. In some implementations, kinesthetic cues are combined with cutaneous cues to improve both contact localization and interpretation of dynamic interaction variables [14, 124].

The main limitation of kinesthetic feedback in wearable prosthetics lies in implementation complexity. Delivering meaningful proprioceptive-like cues often requires forces or displacements larger than those needed for simple cutaneous stimulation, which may increase device size, anchoring requirements, and power consumption [123, 144]. In addition, wearable kinesthetic devices may constrain movement or become less practical for daily use if they require rigid support or large actuation mechanisms. These constraints have limited their translation into compact multi-channel prosthetic feedback

---

[*] For mechanotactile and kinesthetic feedback, the literature is more strongly implementation-dependent than for electrotactile and vibrotactile systems. The reported values should therefore be interpreted as approximate practical ranges rather than strict modality-wide limits.

Table 2.4: Comparison of the main non-invasive feedback techniques used in upper-limb prosthetics.

Criterion	**Electrotactile**	**Vibrotactile**	**Mechanotactile**	**Kinesthetic**
Feedback information	Localized touch cues, distributed tactile events, contact state, force-related coding through stimulation parameters	Contact events, grasp force, slip warnings, state transitions	Pressure, normal contact, distributed pressure, directional skin deformation, force-related tactile cues	Limb motion, force direction, resistance, posture, proprioceptive-like information
Spatial resolution	High	Low to moderate	Moderate	Generally low to moderate
Frequency range	5–100 Hz [141]	50–250 Hz [142]	Typically $\lesssim$ 20 Hz* [138, 143]	Typically $\lesssim$ 10 Hz* [123, 14, 144]
Hardware complexity	Low to moderate	Low	Moderate to high	High
Wearability	High	High	Moderate	Moderate to low
Typical use in prosthetics	Multi-site tactile encoding, sensory substitution, distributed tactile feedback	Event-based cues, force feedback, slip alerts	Pressure rendering, more natural touch substitution, embodiment-oriented designs	Force guidance, proprioceptive augmentation, motion-related feedback
Advantages	High spatial density, compact integration, broad parameter programmability, lightweight hardware	Simple, robust, low cost, easy to implement, intuitive event cueing	More natural contact sensation, intuitive pressure rendering, strong embodiment potential	Addresses proprioceptive dimension, useful for force/motion perception, complements cutaneous feedback
Limitations	Requires calibration, percept varies with skin condition, possible discomfort/habituation	Limited spatial resolution due to actuator size and vibration spread	Mechanical complexity, bulkier integration, limited scalability to many channels	High mechanical complexity, larger actuation demand, lower wearability in compact systems
References	[120, 121, 126, 10, 141]	[120, 121, 145, 146, 10, 142]	[120, 14, 138, 139, 140, 143]	[123, 14, 139, 140, 144]

systems compared with electrotactile and vibrotactile interfaces [120, 10]. Nevertheless, kinesthetic feedback remains highly valuable conceptually because it targets a sensory dimension that is essential for dexterous action but often underrepresented in non-invasive feedback design.

2.6.5 Feedback Modality Selection

Table 2.4 highlights that no single feedback modality is universally optimal; rather, each technique offers a different compromise between perceptual richness, implementation complexity, and suitability for wearable integration. Vibrotactile feedback is attractive for simple and robust cueing, but its limited spatial density makes it less suitable for representing large tactile maps. Mechanotactile and kinesthetic approaches can provide more natural or proprioceptive-like sensations, but their increased mechanical complexity makes scaling to many channels more difficult. Electrotactile feedback, although it requires careful calibration, offers a particularly favorable balance for applications involving distributed tactile sensing because it enables compact multi-electrode layouts, relatively high spatial selectivity, and rapid programmable stimulation updates [126, 147, 148].

For these reasons, electrotactile feedback was selected in this thesis as the most appropriate modality for representing a high-density piezoelectric sensing system. The objective of the proposed human-machine interface is not merely to indicate isolated contact events, but to preserve as much as possible the spatial distribution and the temporal richness of tactile information acquired at the sensing side. In this context, electrotactile stimulation is particularly suitable because it supports dense electrode arrangements over limited skin areas and allows rapid modulation of stimulation parameters in a compact wearable form factor [126, 120, 147, 149]. Thus, the choice of electrotactile feedback is consistent with the central goal of this thesis: transferring distributed tactile information from the artificial skin to the user while maintaining high information density, localized representation, and practical wearability.

2.7 Conclusion

This chapter reviewed the key components required to build tactile-enabled robotic and prosthetic hands, spanning sensing, processing, application-level inference, and feedback delivery. We first outlined the system-level requirements and integration challenges of tactile sensing systems in robotics (e.g., spatial distribution, SNR, crosstalk, wiring/communication, and processing constraints) [44], and summarized typical design targets in terms of force range, spatial resolution, bandwidth, and mechanical/electrical robustness. We then surveyed the main tactile transduction technologies used in e-skin and hand sensorization piezoelectric, resistive, capacitive, and optical highlighting their respective strengths and limitations for hand integration (e.g., bandwidth and dynamic response for piezoelectric sensing [46], static-force capability for capacitive sensors [60], simplicity versus hysteresis/temperature sensitivity for resistive sensors [54, 55], and electromagnetic immunity for optical sensing [24, 23]).

On the data-processing side, we reviewed representative signal-processing strategies used to convert raw tactile time series into compact and discriminative representations. Frequency-domain descriptors derived via FFT can capture contact-induced micro-vibration signatures and have been used successfully for surface characterization and classification [64, 65, 66, 67], but their real-time embedded deployment can be limited by code-size and memory constraints [68]. Time–frequency methods such as DWT provide multiscale representations that can be more expressive when purely spectral features are insufficient [72, 73], while low-cost time-domain statistical descriptors remain attractive when computational efficiency is critical [75, 76]. We also discussed dimensionality reduction and feature-extraction approaches (PCA, ICA, and LDA), which are widely used to compact high-dimensional tactile measurements and/or improve separability [78, 86, 3], with examples spanning material recognition and tactile image processing [79, 90]. Finally, we summarized the role of machine learning for tactile inference, emphasizing that algorithm choice remains problem-dependent (no-free-lunch principle) [91], and that supervised learning continues to dominate prosthesis-oriented tactile perception pipelines [93].

The chapter then surveyed representative tactile applications—hardness and texture classification, as well as grasp stability and slip detection—illustrating that tactile sensing can provide information not directly accessible to vision (e.g., hardness/texture cues and incipient slip) and can support both perception and control. However, the reviewed studies also reveal recurring translation gaps: (i) performance often depends on controlled interaction conditions (e.g., known sliding speed/trajectory for texture features) [74]; (ii) many pipelines are validated offline rather than under real-time constraints where latency, memory footprint, and robustness to non-stationary contact dynamics become limiting factors [68]; and (iii) sensing and inference are frequently studied in isolation from the feedback interface, despite the fact that closed-loop utility ultimately depends on how tactile variables are encoded and perceived.

Finally, we reviewed major non-invasive feedback modalities (vibrotactile, electrotactile, mechanotactile, and kinesthetic) and the encoding parameters that shape the delivered percept (e.g., frequency/amplitude modulation for vibrotactile feedback [125, 126], waveform and electrode/stimulation parameters for electrotactile feedback [133, 10, 134]). Taken together, the literature motivates a system-level approach in which sensing technology, real-time processing, and feedback encoding are co-designed. In particular, the need for scalable and mechanically compliant sensorization, robust low-latency tactile interpretation, and feedback-ready outputs with sufficient bandwidth and repeatability motivates the thesis direction toward PVDF-based distributed sensing and end-to-end, real-time pipelines that can drive structured non-invasive stimulation interfaces for closed-loop prosthetic operation [40, 41, 42].

CHAPTER 3

Experimental Assessment of Tactile Sensing System in a Controlled Environment

3.1 Introduction

Humans can explore diverse objects and simultaneously infer their material properties, a challenging task for robots characterized by the lack of effective sensing systems. This lack can affect different modalities of sensation and control including touch, pressure, and proprioception [150, 151]. Until now, it remains challenging to simulate the human sense of touch as it does not reflect only the tactile response to external physical stimuli (e.g., pressure), but also the material properties (e.g., hardness) through extraction and analysis of information [152]. Hardness, as an important physical property of objects, directly impacts the grasping success of robots. It was proved that discrimination of an object's hardness increases proprioceptive, tactile discriminative, and object recognition senses along with postural control [153]. To this end, hardness perception and intelligent controls can be designed based on the development of tactile sensors [154] and machine learning techniques [155].

Tremendous sensing technologies based on intelligent tactile sensors [152] and machine learning techniques [155] have been developed to recognize the material's properties. Considering hardness as a key feature for efficient motor control of robotic hands [156], many tactile sensors based on various sensing mechanisms were developed and used for hardness recognition. For example, researchers in [157] utilized an optical tactile sensor for tissue hardness discrimination. On the other hand, researchers in [158] developed a membrane-like piezoresistive flexible tactile to discriminate six objects with different hardness. A bionic tactile sensor based on PVDF tactile sensors was developed in [159] and used to detect the hardness of three latex samples during tapping

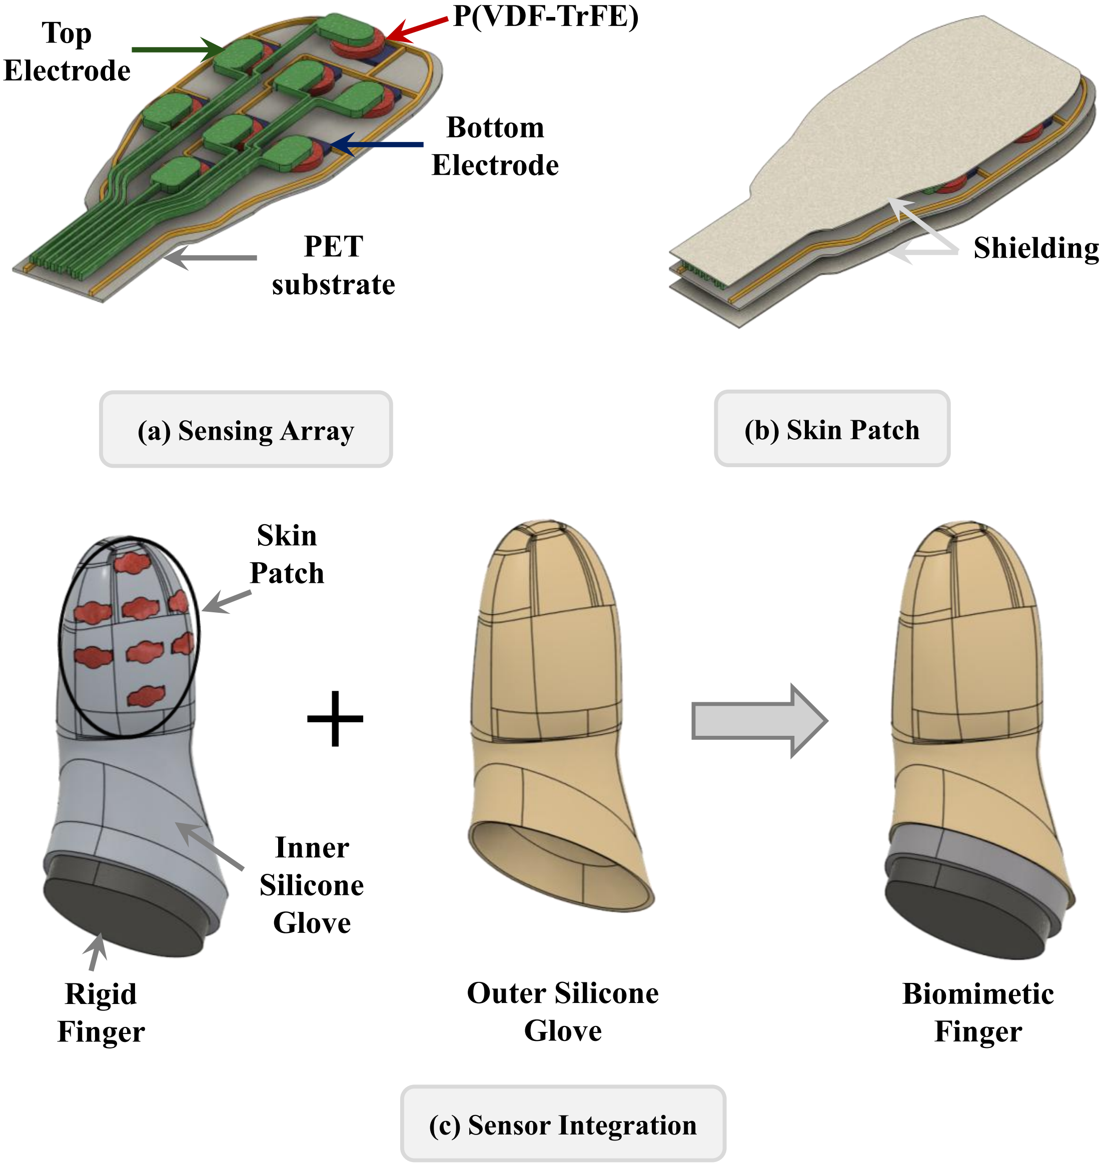

Figure 3.1: (a) structure of the sensor array (b) Sensor Integration

experiments. Neural network algorithms were trained on the time-domain features of the PVDF signals and used to discriminate the hardness of the samples. Authors in [160] calculated the hardness using the load measured by a commercial tactile sensor mounted on an RH8D five-fingered robotic hand. The authors programmed the hand to grasp the object with constant speed and then employed several machine learning models including a support vector machine (SVM) for object recognition. Similarly, authors in

[161] combined a commercial tactile sensor array, principal component analysis (PCA), k-nearest neighbor (KNN), and SVM to recognize the hardness of fruits and vegetables. The authors mounted the sensors on a robotic gripper that manipulates the vegetables at constant force and speed, thus limiting the experimental conditions. Despite the great advancement in hardness classification, most of the works in the literature proposed complex hardness approaches. Moreover, these approaches were validated in a controlled manner in which the gripping speed and load of the manipulator were fixed. This is important as humans use specific movement strategies to extract information about object properties during haptic exploration [162]. Moreover, the relationship between fundamentally different sensory parameters must be known to treat these parameters as relating common object or property [163]. Therefore the load and speed one employs to explore the object affects the perception of its softness since they may alter the intensity and frequency content of the vibrations it elicits in the skin [6]. Therefore, it is crucial to consider the tactile motion alongside the sliding speed and load while developing a tactile sensing system for hardness identification. Moreover, it is generally difficult for robotics to sense the material hardness using one single sensor [164]. To that aim, this study attempted to develop a tactile sensing system that is composed of a PVDF-based biomimetic finger cap, embedded electronics, and Machine Learning (ML) algorithms for hardness discrimination in wearable and soft robots. This chapter demonstrates the system's capability of sensing the relative hardness of objects through a simple operation mechanism. Combined with feature extraction and machine learning, the sensing system can perceive with high accuracy the object's hardness, considering the change in the indentation speed and load. It is organised as follows: Section 3.2 presents the sensing system used in this experiment, briefly illustrating the design and shielding of the sensor, Section 3.3 will present the experimental setup and data collection. This activity was divided into two activities the first one by using the data of one sensor out of eight and it's methodology and results will be presented in Section 3.4, meanwhile the last methodology developed to process the data of eight sensors and it's results will be explained in Section 3.5. Finally, conclusive remarks are given in Section 3.6

3.2 Biomimetic Tactile Sensing System

With the target of implementing the sensing system in different applications, such as robotics (e.g., prosthetic hands) or healthcare (e.g., rehabilitation), the components were integrated into a finger-cap structure that could be mounted on the fingertip (Fig. 3.1b).

The sensing arrays used in this activity were based on a previously developed fully screen-printed flexible sensor technology fabricated by JOANNEUM RESEARCH [165]. Therefore, the focus of the present thesis activity was not on the fabrication of the piezoelectric arrays themselves but on the development of a dedicated integration strategy that embeds these sensing arrays into a biomimetic finger-cap structure suitable for easier usage in applications (e.g., robotics, rehabilitation)

The manufacturing process of the sensing arrays relies on screen printing of ferroelectric sensor arrays based on P(VDF-TrFE) poly(vinylidene fluoride trifluoroethylene) repeated units. Fig. 3.1(a) shows the structure of the sensing arrays. The bottom electrode is screen-printed on a transparent and flexible (175 μm thick) DIN A4 plastic foil (Melinex ST 725) substrate. A ferroelectric polymer P(VDF-TrFE) layer (5.1 μm thick) is then screen-printed onto the bottom electrodes, followed by screen printing of the top electrodes (PEDOT:PSS). A UV-curable lacquer layer is then deposited on top for overall sensor protection. Finally, the poling procedure aligns the dipoles in the thickness direction of the P(VDF-TrFE) crystallites.

In the present work, the sensing arrays were mechanically and electrically prepared for integration into the finger-cap structure. The layout of the sensing arrays and the sensor distribution were chosen to allocate most sensors to the center of the fingertip while allowing practical integration on the finger surface. Before integrating the sensing arrays into the finger cap, protection and shielding layers were added to protect the skin electrically and mechanically while preserving conformity to the finger shape. The sensing array was first connected to a routing PCB that extends the tracks and then sandwiched between two electrically conductive tapes (Model tesa 60262, tesa). The conductive tape was used as a shielding layer to reduce sensitivity to noise. To preserve a high sensor sensitivity, the resulting structure was then embedded between two silicone caps, forming the sensorized finger cap shown in Fig. 3.1. Embedding the sensing array into the silicone caps provided good mechanical contact and enabled efficient transmission of vibrations from the contact surface to the PVDF polymer. The structure of the sensor surface, the intrinsic flexibility of the sensing system, and its wide frequency bandwidth (1 Hz–1 kHz) make it a suitable candidate as a functional constituent of a flexible electronic skin for measuring dynamic contacts.

3.3 Readout Circuit

To acquire the sensor's response, we utilized the embedded electronics developed and validated in [27] (Fig. 3.1b)). It is composed of an ultra-low-power microcontroller based on an ARM Cortex M0 (BL600 module, Laird Connectivity, US), a DDC232 (Texas Instrument, US) current-input analog-to-digital converter, and a USB data transfer interface (i.e. FTDI232). The design can handle up to 32 sensors through two sockets with 16 input channels. Each channel is connected to an offset circuit to adjust the baseline of the bipolar signals generated by the sensors, which allows DDC232 to receive both the positive and negative polarities of the sensor signals. The DDC232 is configured to use 16-bit resolution and cover the maximum input charge response. In this study, the embedded electronics was configured to simultaneously sample the 32 input channels at a sampling rate of up to 2 kHz. The electronics features data collection and transmission through a USB connection.

3.4 Experimental Setup and Data Collection

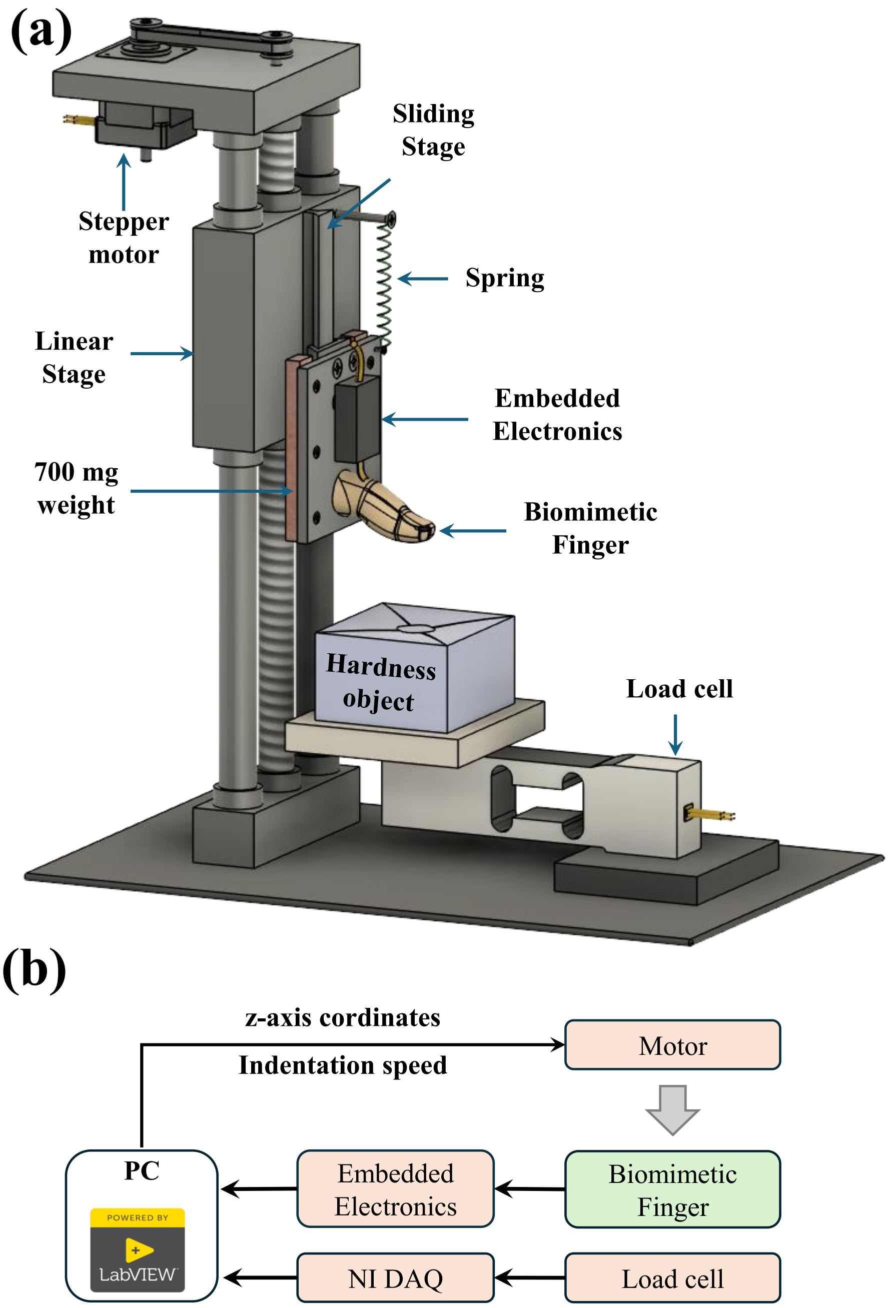

Figure 3.2: Sketch of the experimental setup

To collect experimental sensor data for hardness discrimination, 3D-printed objects presented and used in [166] were utilized in this study. Five cubic objects (7 x 7 x 4 cm) were 3D printed with five different filling percentages (3, 5, 7, 10, and 12%) using thermoplastic polyurethane (TPU). Figure 3.1c) shows the experimental setup used to collect the experimental sensor data. A 3-axis cartesian robot was used to perform repeated indentation actions on the 3D-printed objects while varying the indentation speed and load. In this study, the z-axis was used to indent the hardness object. The cartesian robot is composed of two moving parts: the main part is a thick aluminum plate (plate 1) mounted on two parallel linear stages that were oriented vertically to form the z-axis. The linear stage (i.e., the main part) is driven by a stepper motor and is controlled by modulating the z-coordinates and indentation speed. A second aluminum plate (plate2) was attached to a slider on the main part and hung through a spring as shown in Figure 3.2.

A 700 mg weight was mounted on the second plate to modulate the load on the surface of the skin patch gradually. The finger cap and embedded electronics were mounted on a 3D finger-shaped support and then fixed to plate2 (Fig. 3.1). To measure the load applied through the z-axis on the objects, the 3D objects were placed on a load cell (Tedea Huntleih, Model 1042) mounted in line with the z-axis as shown in Fig. 3.2). In particular, the load cell was used to set the z-coordinates that result in different indentation loads. The load cell response was processed and conditioned using a DAQ (NI, US). Data acquisition, visualization, collection, and control of the Cartesian robot were implemented using a graphical user interface (GUI) developed with NI LabVIEW on a host PC.

In this study, indentation experiments were planned to perform human-like contact to perceive the hardness of the objects in real-time while controlling the speed and load. To this end, a programmable mechanical input was applied by programming the Cartesian to indent the 3D objects using the sensorized finger while continuously measuring the electrical response of the PVDF sensors and the reference load cell. The indentation procedure consisted of the following three steps at a programmed Indentation Speed (IS): Initial indentation (initial load: 8 N), increasing the load until reaching the Target Loads (TL), decreasing the load to the initial load, and then releasing the finger gradually. Figure 3.3 shows an example of the indentation sequence at two combinations of IS and TL and the corresponding eight PVDF sensor outputs.

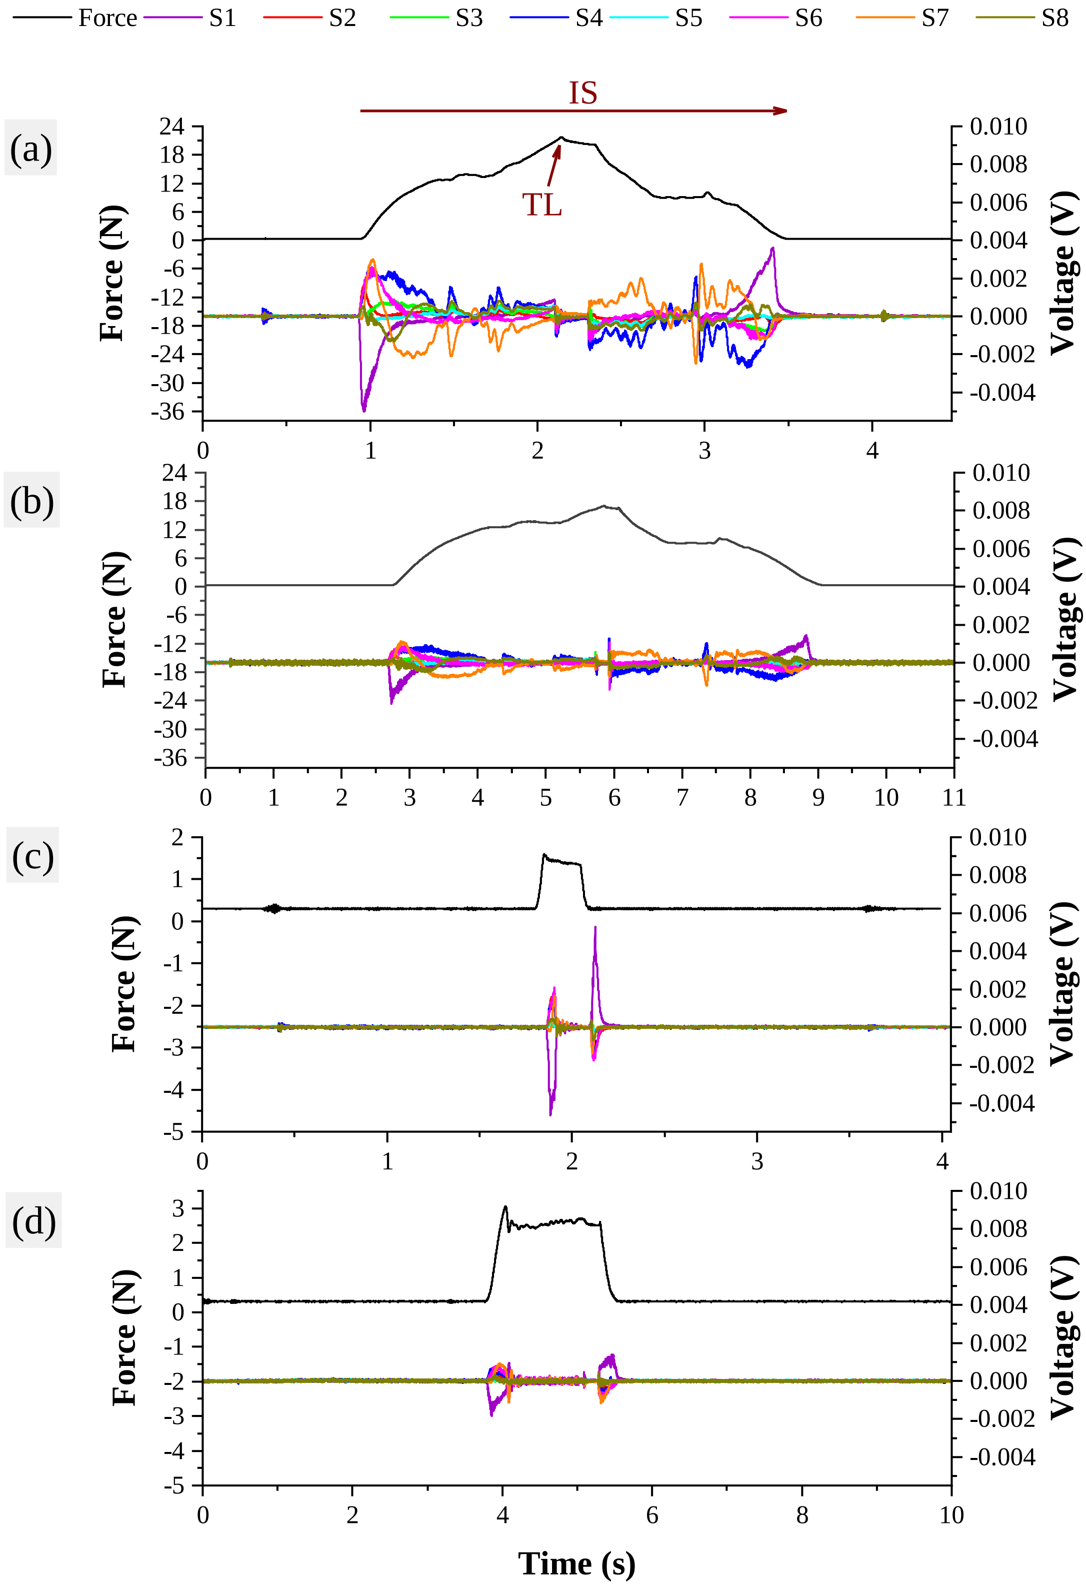

Figure 3.3: Example of the tactile signal from the 8 sensors in response to the indentation at (a) TL = 11 N; IS = 7.2 mm/s, (b) TL = 11 N; IS = 2.3 mm/s, (c) TL = 8.5 N; IS = 7.2 mm/s, (d) TL = 8.5 N; IS = 2.4 mm/s.

Notably, the indentation stimulus measured by the load cell is not aligned with the sequence programmed and applied by the cartesian robot. For example, a rapid decrease is detected in the indentation stimulus around 3 seconds, this is identified as a spontaneous event due to the spring used to modulate the load gradually. The indentation procedure was repeated for 4 TL {8.5, 9.5, 10.5, 11} N at 3 IS {2.4, 5.6, and 7.2} (mm/s). Then repeated for each 3D-printed object, thus generating a dataset composed of 60 combinations of objects, TL, and IS. Finally, each combination was repeated for 10 trials. The dataset can be formulated as $\mathcal{D} = \left\{(X_i, y_i)\,,\ X_i \in \mathbb{R}^{N_C \times N_S},\ y_i \in \{five\ hardnesses\}\,;\ i = 1, ..., 600\right\}$, where X_i is a 2D tensor with $N_C = 8$ representing the number of channels (i.e., the sensors) and N_S the number of samples that depends on the TL and IS as shown in Figure 3.3

3.5 Activity 1: Hardness Classification using one sensor

3.5.1 Methodology

To reduce the complexity of the proposed approaches, we consider selecting only the most active sensors. For that reason, the energy along each sensor was computed for each $X_i \in \mathcal{D}$ to determine the most active sensor. As a result, a new dataset $\tilde{\mathcal{D}}$ was extracted from $\mathcal{D}$, which can be formulated as $\tilde{\mathcal{D}} = \left\{\left(\tilde{X}_i, y_i\right)\,,\ \tilde{X}_i \in \mathbb{R}^{1 \times N_S},\ y_i \in \{...\}\right\}$. Inspired by the work done in [161] and [167], five features were extracted for each $\tilde{X}_i \in \tilde{\mathcal{D}}$. Specifically, the set of features $f5 = \{$Standard Deviation (STD), Skewness (SK), Kurtosis (KT), Peak-to-Peak (PP), STD of differences between adjacent samples (SD)$\}$. Consequently, each $\tilde{X}_i \in \tilde{\mathcal{D}}$ was transformed to $\tilde{X}_i^{f5} \in \tilde{\mathcal{D}}^{f5}$, where $\tilde{\mathcal{D}}^{f5} = \{(\tilde{X}_i^{f5}, y_i),\ \tilde{X}_i^{f5} \in \mathbb{R}^{1 \times 5},\ y_i \in \{...\}\}$. The principal component analysis (PCA) technique [168] was then applied on $\tilde{\mathcal{D}}^{f5}$ to determine the best representative feature, thus reducing the dataset dimensions. Figure 3.4 shows the PCA biplot computed on $\tilde{\mathcal{D}}^{f5}$ to insight the best combination of features. The length of the feature vectors and the angles between the lines reflect the relevance and the correlation of the represented features [169]. Hence, features were ranked by multiplying the biplot coordinate values of each feature by the corresponding component explainability, ranking the extracted features as follows $\{STD \rightarrow PP \rightarrow SD \rightarrow SK \rightarrow KT\}$. To study the impact of the number of features (i.e., computation complexity), three subsets of features, $f2$, $f3$, and $f4$,

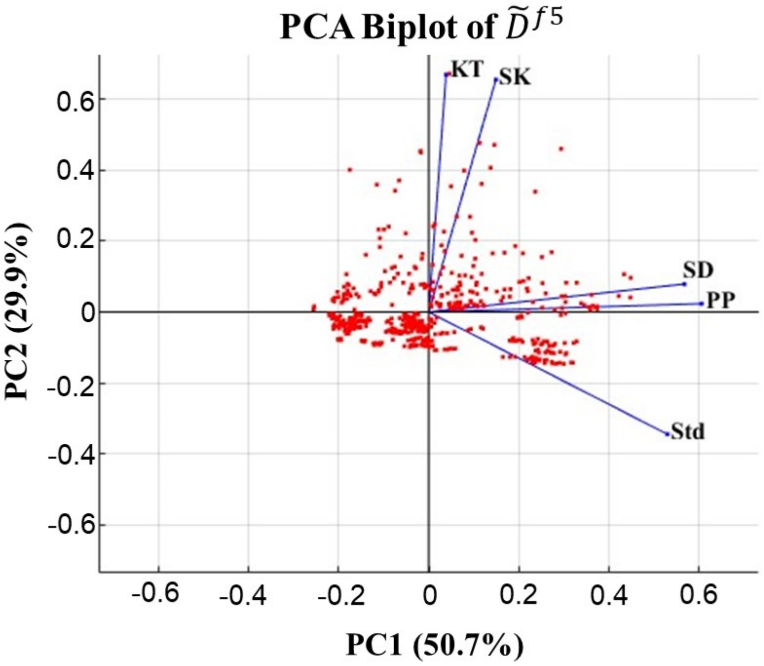

Figure 3.4: PCA Biplot of the extracted features in $\tilde{\mathfrak{D}}^{f5}$. Blue lines represent the feature vectors indicating the relevance of each feature

were selected from the list of the ranked features. The list subsets $f2$, $f3$, and $f4$ represent the most two, three, and four relevant features respectively. As a result, from each $\tilde{X}_i^{f5} \in \tilde{\mathfrak{D}}^{f5}$, $\tilde{X}_i^{f2}$, $\tilde{X}_i^{f3}$, and $\tilde{X}_i^{f4}$ were extracted, with f2 = {STD, PP}, f3 = {STD, PP, SD}, and f4 = {STD, PP, SD, SK}. Hence, three more datasets were generated, i.e., $\tilde{\mathfrak{D}}^{fj} = \{(\tilde{X}_i^{fj}, y_i),\ \tilde{X}_i^{fj} \in \mathbb{R}^{1\times j},\ y_i \in \{...\}\}$ with $j = \{2,3,4\}$ indexing the subset of features. The four datasets $\tilde{\mathfrak{D}}^{f2}$, $\tilde{\mathfrak{D}}^{f3}$, $\tilde{\mathfrak{D}}^{f4}$ and $\tilde{\mathfrak{D}}^{f5}$ were used to train and test the ML algorithms described in the next section.

To demonstrate the impact of IS and TL on the discrimination capability of the developed system, each of datasets $\tilde{\mathfrak{D}}^{f2}$, and $\tilde{\mathfrak{D}}^{f5}$ were divided into 12 sub-datasets representing the combination between the IS and TL. As a result, each $\tilde{X}_i^{f2} \in \tilde{\mathfrak{D}}^{f2}$ and $\tilde{X}_i^{f5} \in \tilde{\mathfrak{D}}^{f5}$, was transformed into $\tilde{X}_{ic}^{f2} \in \tilde{\mathfrak{D}}_c^{f2}$ and $\tilde{X}_{ic}^{f5} \in \tilde{\mathfrak{D}}_c^{f5}$ respectively. Where $\tilde{\mathfrak{D}}_c^{f2} = \{(\tilde{X}_{ic}^{f2}, y_i),\ \tilde{X}_{ic}^{f2} \in \mathbb{R}^{1\times k},\ y_i \in \{...\};\ i = 1, ..., 50\}$, $\tilde{\mathfrak{D}}_c^{f5} = \{(\tilde{X}_{ic}^{f5}, y_i), \tilde{X}_{ic}^{f5} \in \mathbb{R}^{1\times c},\ y_i \in \{...\};\ i = 1, \dots, 50\}$ and c representing the IS and TL combinations $\{[TL = 8.5,\ IS = 2.4], \dots, [TL = 11,\ IS = 7.2]\}$. The 12 datasets $\tilde{\mathfrak{D}}_c^{f2}$ ($c = \{1, \dots, 12\}$) were also used to train and test the ML algorithms.

3.5.2 Hardness discrimination algorithms

Three discrimination algorithms were evaluated as they met with significant success in numerous real-world discrimination tasks [167, 161, 170]. The algorithms are a support vector machine (SVM), a Single-layer Feed-Forward Neural Network (SLFNN), and a K-nearest neighbor (KNN).

3.5.2.1 Support vector machine (SVM)

The SVM classifier is a supervised ML algorithm that computes the hyperplane that maximizes the margin to the nearest samples of the classes. The Gaussian kernel was employed to project the data into a new space.

3.5.2.2 Single-layer Feed-Forward Neural Network (SLFNN)

An SLFNN is a fully connected network with only one single hidden layer trained through the gradient-descent backpropagation technique. The ReLU function was used as an activation function for the hidden neurons, and an output layer comprised 5 neurons to classify the 5 hardness classes.

3.5.2.3 K-Nearest Neighbor (KNN)

The class of a test sample is determined by taking the majority vote of class labels among the K-closest training samples. The Euclidean distance was employed.

All the algorithms were trained and tested on 5 folds using a stratified cross-validation technique applied to each dataset, where the stratification balances the number of samples per class in each fold. Then, 4 out of 5 folds, in rotation, were used for training and the remaining fold was used for the testing. The training set was subsequently divided into training (80%) and validation (20%) sets. The validation set was employed to determine the optimal configuration of hyperparameters. Specifically, for the SVM, this included the standard deviation for the Gaussian kernel and the l2-regularizer parameter; for the SLFNN, it encompassed the number of neurons in the hidden layer and the l2-regularizer parameter; while for the KNN, optimization involved the K parameter.

Table 3.1: Average discrimination accuracy (Acc) and standard deviation (SD) of the ML algorithms over 5 Folds on the four datasets.

	Acc ± SD (%)		
Dataset	SVM	SLFNN	KNN
$\tilde{\mathscr{D}}^{f2}$	82.2 ± 3	77.3 ± 6	80.7 ± 4
$\tilde{\mathscr{D}}^{f3}$	89.3 ± 2	86.3 ± 9	88.3 ± 1
$\tilde{\mathscr{D}}^{f4}$	91.8 ± 4	91.8 ± 8	92.3 ± 2
$\tilde{\mathscr{D}}^{f5}$	96.5 ± 2	96.2 ± 6	94.6 ± 1

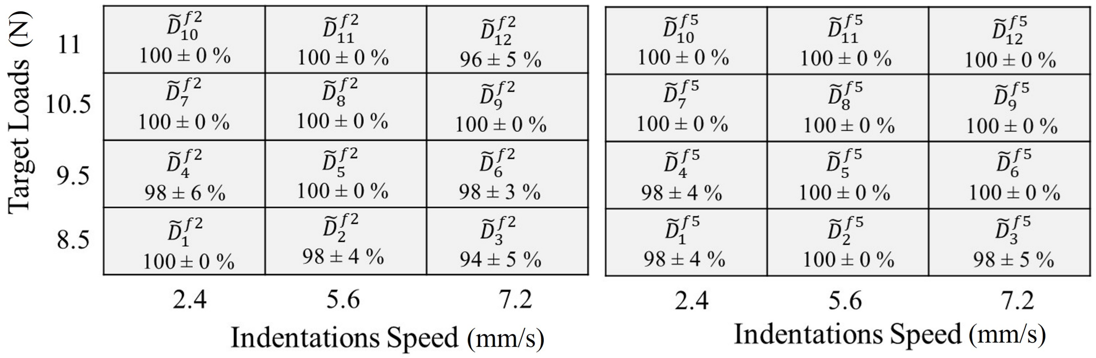

Figure 3.5: Average discrimination accuracy of SVM algorithm when trained on datasets with fixed IS and TL.

3.5.3 Results and Discussion

Table 3.1 summarizes the average discrimination accuracy for each ML algorithm over four datasets (section 3.5). The three models exhibit an increase in discrimination accuracy as a function of the number of features, thus highlighting the importance of the extracted features. In general, The three ML algorithms performed similarly over the four datasets, however, SLFNN possesses the lowest accuracy of 77.3 % when trained on the smallest dataset $\tilde{\mathscr{D}}^{f2}$ (2 features out of 5). As expected, training the ML algorithm using $\tilde{\mathscr{D}}^{f2}$ results in an increase of $\gtrsim$ 10 % in the average discrimination accuracy for all the algorithms. Therefore involving the SD features ($\tilde{\mathscr{D}}^{f3}$) in training the three algorithms results in a $\sim$ 10 % increase in the accuracy confirming the outcome of the PCA biplot. This is confirmed when adding less important features to the training dataset ($\tilde{\mathscr{D}}^{f4}$), in which a lower increase in the accuracy was noticed for the three models (91.8 ± 4 % for SVM and SLFNN and 92.3 ± 2 % for KNN). When trained on $\tilde{\mathcal{D}}^{f5}$, SLFNN and SVM performed similarly and achieved the highest accuracies of (96.5 ± 2 %) and (96.2 ± 6 %) respectively. Similarly, KNN achieved a satisfactory accuracy (94.6 ± 1 %) when trained on the same dataset, however with a lower standard deviation. It is important to note that KNN and SVM have outperformed the SLFNN algorithm over the four datasets in terms of the standard deviation thus showing a stable performance over the 5 folds and generalization. As a result, SVM outperforms SLFNN and KNN for all the datasets in which it exhibits the highest discrimination accuracy (96.5 ± 2 %) and low standard deviation when trained on dataset $\tilde{\mathcal{D}}^{f5}$ which contains the features extracted from the whole dataset ($\tilde{\mathcal{D}}$) i.e, varying IS and TL. For that reason, the SVM algorithm was used to demonstrate the impact of varying IS and TL on its discrimination behavior. Figure 3.5 illustrates discrimination behavior of the SVM algorithms when trained on $\tilde{\mathcal{D}}_c^{f2}$ and $\tilde{\mathcal{D}}_c^{f5}$ datasets.

This study demonstrates the capability of the tactile sensing system to discriminate five different levels of object hardness using feature extraction and ML algorithms. Similar to the work done by [167], the proposed discrimination approach achieves an optimal discrimination accuracy of $\sim$ 100% when trained on one of the datasets characterized by fixed IS and TL ($\tilde{\mathcal{D}}_c^{f2}$) and the response of a single sensor. Moreover, the proposed approach achieved an accuracy of 96.5 ± 2 % using only 5 features extracted from the response of a single sensor (16 sensors in [167]) and at variable IS and TL. The SVM algorithm trained on the dataset $\tilde{\mathcal{D}}^{f5}$ achieved an accuracy higher than that in [169] using the same algorithm. Compared to the discrimination approaches presented in [169] and [161], the proposed approach employed PCA to rank the extracted features instead of using it to extract features and train the ML algorithm. This results in a satisfactory accuracy using only 2 features from a single sensor, thus reducing the complexity of the proposed approach. So far, the available approaches haven't considered speed and load as key features for the discrimination problem. However, the current study proved that varying these parameters influences the discrimination behavior of the proposed sensing system. Another important aspect offered by the proposed approach is the scalability in which its complexity remains acceptable when the system is scaled up in terms of input matrix size and number of classes. In terms of generalization, the proposed approach can be applied to solve new problems (e.g. texture) by fine-tuning or retraining the ML algorithms with new data. Thus, the system must be validated in more active scenarios where daily-life objects are grasped at random loads and speeds.

3.6 Activity 2: Real-Time Simulation of Hardness Classification using All Sensors

3.6.1 Methodology

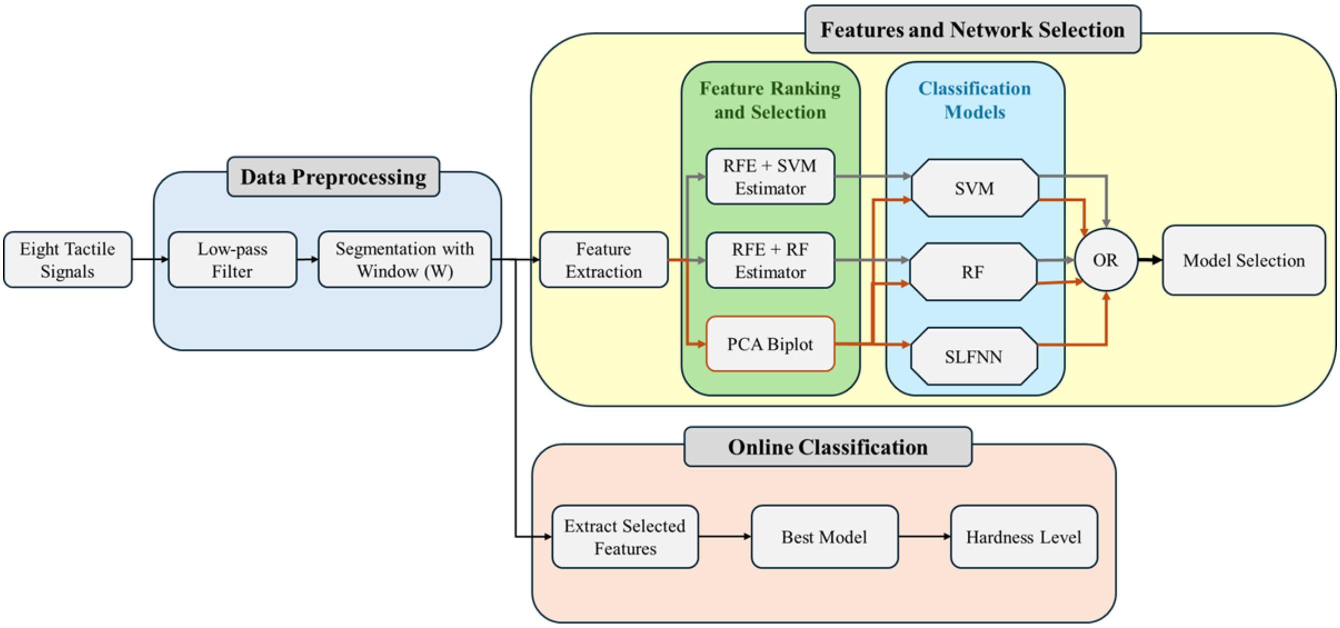

Figure 3.6: Block diagram for piezoelectric data processing

The objective of this study is to achieve hardness classification through a structured machine learning framework. Since human tactile perception relies on continuous sensory feedback rather than batch processing, the proposed approach adopts an incremental data-feeding strategy to the elaboration pipeline, mimicking human perception.

As illustrated in the block diagram in Figure 3.6, the data were first filtered and segmented to generate new datasets, which were then processed in the "Features and model selection" block. In this stage, the most relevant features were selected and used to train different ML algorithms. The optimal combination of features and ML model was subsequently employed in the “Online classification” block. In a real-world implementation, the filtered and segmented data can be continuously fed into this block to classify different hardness levels in real time.

3.6.1.1 Signal filtering

A low-pass filter with a cutoff frequency of 30 Hz was applied to $X_i \in \mathcal{D}$ to remove high-frequency noise while preserving relevant tactile signal components, as described in [171]. Moreover, the parts of the signals not related to the indentation action—specifically, the portions before the touch onset and after the touch release in Figure 3.7 were removed. The resulting filtered data are denoted as $\widehat{X_i} \in \mathbb{R}^{N_C \times \widehat{N_S}}$ in the following, where $\widehat{N_S} < N_S$ due the removal of non-touch-related portions of the signal.

3.6.1.2 Segmentation

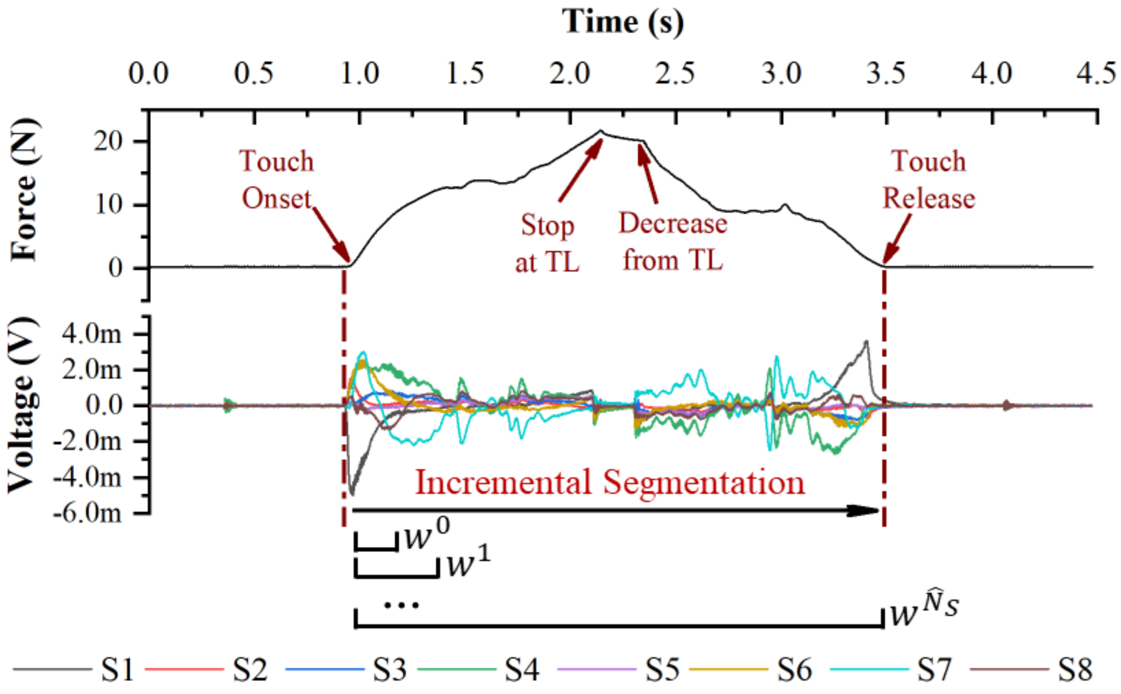

Figure 3.7: The figure shows the incremental segmentation method adapted to create $\mathcal{D}_p$ and $\mathcal{D}_{100}$

The filtered signals were then divided into windows for the steps. In a real-time implementation, the online classification block is continuously fed with incremental windows to capture fine temporal variations while retaining information from previous segments. Accordingly, each filtered signal $\widehat{X_i}$ was segmented into a sequence of windows $W^1, W^2, \ldots, W^{\widehat{N_S}}$. The first window W^1 contains the initial $n < \widehat{N_S}$ samples of $\hat{X}_i$, whereas W^2 include W^1 concatenated with an additional window of n samples. The rationale for the construction of the subsequent windows follows the same principle and is illustrated in Figure 3.7.

Two segmentation strategies were employed:

1. *Incremental fixed-size windowing.* In this approach, an initial window of $n = 100$ samples, denoted as W_i^{100}, was selected. Subsequently, the window size was incrementally increased (200, 300, and so on) until the entire signal was encompassed. The dataset resulting from this segmentation strategy is defined as:

$$\mathcal{D}_{100} = \left\{ \left(W_i^k, y_i\right) \middle| \begin{array}{l} W_i^k \in \mathbb{R}^{N_C \times k}, \\ y_i \in \{\text{five classes}\}; \\ i = 1, \ldots, N \end{array} \right\} \tag{3.1}$$

   where k represents the incremental window sizes, $k \in \{100, 200, 300, ..., \widehat{N_S}\}$. Because indentation trials differ in duration due to variations in TL and IS values, the number of windows per trial also varies: longer trials yield more windows, whereas shorter ones produce fewer.

   This segmentation strategy can be applied in real-world scenarios by implementing a thresholding mechanism to detect touch onset and release, thereby triggering the classification process only when relevant information is available.

2. *Progressive percentage windowing.* To investigate how the ML models behave across different sections of the signal (e.g., around touch onset and touch release) and how hardness-related information is distributed within them, the signals are segmented into windows corresponding to specific percentages of $\widehat{X_i}$. Specifically, W^{10} contains the first 10% of the signal, W^{20} the 20%, and so forth.

This dataset is formulated as follows:

$$\mathcal{D}_p = \left\{ (W_i^p, y_i) \,\middle|\, \begin{array}{l} W_i^p \in \mathbb{R}^{N_C \times p}, \\ y_i \in \{\text{five classes}\}, \\ i = 1, \ldots, 6000 \end{array} \right\} \tag{3.2}$$

where $p \in \{10\%, \ldots, 100\%\}$.

It is worth noting that this segmentation technique cannot be applied in real scenarios where the signal length is not known a priori. It is employed solely for comparison purposes with the incremental fixed-size windowing approach.

3.6.2 Feature Extraction and Selection

To reduce the number of features and retain those that best represent the tactile dynamics, a feature extraction and selection pipeline is applied, as illustrated in Fig. 3.6. For each windowed trial $W \in \mathbb{R}^{N_C \times L}$, where $N_C = 8$ is the number of sensing channels and L is the number of samples in the considered window, let $w_c[t]$ denote the t-th sample of channel c, with

$$c = 1, \ldots, N_C, \qquad t = 1, \ldots, L.$$

For each channel c, six time-domain descriptors are extracted. Let

$$\mu_c = \frac{1}{L} \sum_{t=1}^{L} w_c[t]$$

denote the sample mean of channel c. The extracted features are then defined as follows:

Standard deviation

$$\text{StD}_c = \sqrt{\frac{1}{L} \sum_{t=1}^{L} \left(w_c[t] - \mu_c \right)^2}.$$

This feature measures the dispersion of the signal around its mean and reflects the variability of the tactile response within the considered window.

Skewness

$$\mathrm{Skewness}_c = \frac{1}{L}\sum_{t=1}^{L}\left(\frac{w_c[t]-\mu_c}{\mathrm{StD}_c}\right)^3 .$$

Skewness quantifies the asymmetry of the signal distribution with respect to its mean.

Kurtosis

$$\mathrm{Kurtosis}_c = \frac{1}{L}\sum_{t=1}^{L}\left(\frac{w_c[t]-\mu_c}{\mathrm{StD}_c}\right)^4 .$$

Kurtosis measures the peakedness of the signal distribution and is sensitive to impulsive or sharply varying responses.

Maximum value

$$\mathrm{Maximum}_c = \max_{1\le t\le L} w_c[t].$$

This feature captures the largest signal amplitude observed in the window.

Minimum value

$$\mathrm{Minimum}_c = \min_{1\le t\le L} w_c[t].$$

This feature captures the smallest signal amplitude observed in the window.

Peak-to-peak amplitude

$$\mathrm{PtP}_c = \mathrm{Maximum}_c - \mathrm{Minimum}_c.$$

Peak-to-peak amplitude describes the dynamic range of the signal over the considered window.

The six-dimensional feature vector associated with channel c is therefore defined as

$$\boldsymbol{\phi}_c(W) = \begin{bmatrix}\mathrm{Skewness}_c & \mathrm{Kurtosis}_c & \mathrm{StD}_c & \mathrm{PtP}_c & \mathrm{Maximum}_c & \mathrm{Minimum}_c\end{bmatrix}^\top \in \mathbb{R}^6.$$

The complete feature vector associated with the window W is then obtained by concatenating the descriptors extracted from all channels:

$$z(W) = \begin{bmatrix}\boldsymbol{\phi}_1(W)^\top & \boldsymbol{\phi}_2(W)^\top & \cdots & \boldsymbol{\phi}_{N_C}(W)^\top\end{bmatrix}^\top \in \mathbb{R}^{6N_C}.$$

Since $N_C = 8$, each window is represented by a 48-dimensional feature vector. These descriptors were selected because they are low-cost to compute and can be updated incrementally from local signal statistics, which makes them compatible with online or memory-constrained implementations.

This dimensionality is subsequently reduced through three ranking and selection strategies, as illustrated in the "Feature ranking and Selection" block in Fig. 3.6. The first strategy is based on principal component analysis (PCA), while the other two rely on recursive feature elimination (RFE) with different estimators.

3.6.2.1 PCA

PCA [172] is an unsupervised dimensionality reduction technique that transforms the original correlated features into a new set of orthogonal components, known as principal components (PCs). These components are linear combinations of the original variables and are ordered according to the amount of variance they explain in the dataset. By projecting the data into a lower-dimensional subspace defined by the first few components (two or three), PCA helps identify the most influential features and visualize their relationships through a biplot representation, where features with higher loadings are considered more relevant.

Similar to the previous work reported in [173], PCA is used here as a feature-ranking tool. Let $\ell_{j,1}$ and $\ell_{j,2}$ denote the loadings of feature j on the first and second principal components, respectively, and let η_1 and η_2 be the corresponding explained-variance ratios. An importance score is assigned to each feature according to

$$s_j = \eta_1 |\ell_{j,1}| + \eta_2 |\ell_{j,2}|.$$

Features are then ranked in descending order of s_j, and the first 10 elements of the resulting ordered list define the PCA-based subset used for model training. In this way, PCA provides an unsupervised criterion that emphasizes descriptors aligned with the dominant directions of variance in the feature space.

For both segmentation strategies, the PCA-based ranking was recomputed independently within each outer training fold, and the resulting top-10 subset was then applied unchanged to the corresponding held-out fold.

3.6.2.2 RFE

The RFE [174] is a supervised feature selection method that iteratively trains a model and removes the least significant features to identify the most relevant subset. Starting from the full feature set, an estimator is fitted to the data, feature importance is computed, and the least important features are recursively discarded until a predefined number remains. This process produces a ranked list reflecting each feature's contribution to the model's predictive performance, thereby improving generalization and reducing redundancy among correlated features.

In this study, RFE is implemented with two estimators, namely SVM and RF. Let

$$\mathcal{F} = \{f_1, f_2, \ldots, f_{48}\}$$

denote the set of features extracted from a window. For a given estimator $m \in \{\mathrm{SVM}, \mathrm{RF}\}$, RFE produces an ordered ranking

$$\mathcal{F}_{\mathrm{RFE}}^{(m)} = \big(f_{\pi^{(m)}(1)}, f_{\pi^{(m)}(2)}, \ldots, f_{\pi^{(m)}(48)}\big),$$

where $\pi^{(m)}(\cdot)$ is a permutation that orders the features from most to least relevant according to estimator m. For the SVM-based configuration, feature relevance is derived from the absolute magnitude of the model coefficients, whereas for the RF-based configuration it is obtained from impurity-based importance measures. In both cases, the first 10 elements of $\mathcal{F}_{\mathrm{RFE}}^{(m)}$ define the subset used to train and tune the corresponding classifier. Since RFE is not used with the SLFNN in this work, no RFE-based subset is defined for that model.

Likewise, for both segmentation strategies, the RFE ranking was recomputed independently within each outer training fold using the corresponding estimator, and the resulting top-10 subset was transferred unchanged to the corresponding held-out fold.

3.6.3 Model Selection

As previously described, an incremental classification approach is adopted, where the classifier receives progressively larger temporal segments of the signal. Therefore, the selected models must efficiently process incremental updates while remaining lightweight and suitable for real-time or resource-constrained deployment.

Although recurrent neural networks, such as Long Short-Term Memory (LSTM) models, are well suited for sequential data and can capture temporal dependencies, their deployment on embedded systems is often impractical due to considerable computational and memory demands—particularly in online scenarios requiring low-latency classification [175][176].

Conversely, one-dimensional convolutional neural networks (1D-CNNs) effectively capture local temporal patterns within fixed-length windows of time-series data but are less suited to the incremental classification paradigm. Standard 1D-CNNs require fixed-length inputs, forcing shorter early windows to be padded, introducing artifacts or bias or longer ones to be truncated or interpolated, thereby losing information [177]. Moreover, CNNs process each window independently and lack an inherent mechanism to retain knowledge from previous segments, which is critical in incremental learning where new windows build upon past information. Although memory mechanisms such as hybrid CNN–RNN architectures, temporal attention, or external buffers can be integrated, these techniques increase architectural complexity and reduce interpretability.

For the aforementioned reasons, we adopt lighter models—SVM, RF, and SLFNN—that better align with incremental update strategies and deployment constraints. These models enable efficient retraining or adaptation with minimal computational cost and offer greater interpretability and robustness, particularly in settings with limited datasets and structured feature representations.

3.6.3.1 Support Vector Machine (SVM)

Two kernel types are evaluated, linear and radial basis function (RBF)—to capture both linear and nonlinear relationships. For the RBF kernel, the kernel coefficient γ is tested using the `scale` and `auto` options available in the Python library *scikit-learn*, allowing the kernel width to adapt to the feature distribution. A one-vs-one training strategy is employed for multi-class classification.

3.6.3.2 Random Forest (RF)

To optimize model performance, several hyperparameters are tuned. The number of trees in the ensemble is tested at 100 and 200, while the maximum tree depth is fixed at 5 to limit model complexity. The minimum number of samples required to split an internal node is evaluated at 2 and 5, and the minimum number of samples required at a leaf node is set to either 1 or 2.

3.6.3.3 Single Layer Feedforward Neural Network (SLFNN)

The SLFNN is a lightweight neural architecture composed of a single hidden layer with nonlinear activation. During model tuning, the hidden layer size is tested with 50 and 100 neurons, while the activation function is fixed to ReLU to introduce nonlinearity. The L2 regularization term is varied between 0.0001 and 0.001 to mitigate overfitting, and two learning rate schedules—constant and adaptive—were compared to evaluate their effects on convergence stability and speed.

3.6.3.4 Hyperparameter tuning and model selection protocol

For each of the two segmentation strategies, namely $\mathcal{D}_p$ and $\mathcal{D}_{100}$, the same model-selection protocol was applied. Five candidate pipelines were considered:

$$\mathrm{RF}_{\mathrm{PCA}}, \quad \mathrm{SLFNN}_{\mathrm{PCA}}, \quad \mathrm{SVM}_{\mathrm{PCA}}, \quad \mathrm{SVM}_{\text{RFE-SVM}}, \quad \mathrm{RF}_{\text{RFE-RF}}.$$

Within each segmentation dataset, the same five stratified outer folds were used for all candidate pipelines in order to ensure a fair comparison. Stratification was performed over the combinations of hardness level, target load, and indentation speed, so that each fold preserved the distribution of the experimental scenarios.

For each outer fold, the complete feature-ranking stage was recomputed using only the corresponding training partition. In the PCA-based pipelines, the PCA ranking was estimated exclusively on the training fold and the top 10 ranked features were retained. In the RFE-based pipelines, the RFE ranking was likewise estimated exclusively on the training fold using the corresponding estimator, and again the top 10 ranked features were retained. The resulting selected subset was then applied unchanged to the corresponding held-out fold.

Within the training partition of each outer fold, hyperparameter tuning was performed by grid search using a three-fold stratified inner cross-validation procedure. For every candidate hyperparameter configuration, the model was trained and validated only on the training portion of that fold, and the configuration achieving the highest mean inner validation accuracy was selected. The selected configuration was then refitted on the full outer-training partition and evaluated on the corresponding held-out fold. Repeating this procedure over the five outer folds yielded both fold-wise accuracies and out-of-fold predictions for each candidate pipeline.

For each segmentation strategy, the final ranking of the candidate pipelines was based on the mean accuracy across the five outer folds. The corresponding out-of-fold predictions were also retained in order to support the paired statistical analyses reported later in the chapter [178, 179].

3.6.4 Online classification

The final evaluation phase was conducted to assess model performance under a realistic testing setup that emulates online operation. The original dataset $\mathcal{D}$ was divided into two disjoint subsets for training and testing. For each hardness level and each combination of force and speed, ten trials were available; seven trials were allocated for training and the remaining three for testing. The same windowing strategies described in Sec. 3.6.1.2 were independently applied to both the training and testing sets. To ensure a balanced training set, a small number of windows were randomly removed from overrepresented classes, ensuring that all hardness, target load, and indentation speed levels contained an equal number of samples.

For each segmentation strategy, the best-performing pipeline identified during the model-selection stage was subsequently retrained on the training portion of the original dataset and evaluated on the corresponding held-out test trials. The train/test evaluation was therefore kept strictly separate from the five-fold cross-validation procedure used for model selection.

More specifically, for each segmentation dataset, the same feature-selection strategy retained during model selection was fitted using only the training trials and then applied unchanged to the test trials. The selected model, together with the hyperparameter configuration trained during model selection, was then retrained on the training windows and evaluated on the held-out test windows. For each test trial, the model output was stored together with the true label, enabling both quantitative and qualitative analysis of the final selected classifier under previously unseen interactions [178, 179].

3.6.5 Statistical analysis

For this activity, formal inferential analysis was performed only on the fixed-step cumulative windowing results, since the percentage-based windowing strategy was used only for descriptive analysis of information emergence and not for the final model-selection procedure. Five candidate pipelines were considered in the fixed-step setting: RF trained on the top 10 PCA-ranked features, SLFNN trained on the top 10 PCA-ranked features, SVM trained on the top 10 PCA-ranked features, SVM trained on 10 features selected by RFE with a linear SVM estimator, and RF trained on 10 features selected by RFE with a Random Forest estimator.

Given that model selection was based on five matched outer cross-validation folds, the fold-wise accuracies of the five pipelines were first compared using the Friedman test, which is a standard non-parametric choice for comparing more than two matched classifiers across repeated evaluation blocks [180]. Since only five outer folds were available, fold-level post-hoc testing was expected to have limited statistical power. Therefore, to complement the fold-level analysis, the out-of-fold predictions produced by the five pipelines on the same held-out samples were also compared at the sample level. In this second analysis, Cochran's Q test was used as an omnibus test for matched binary outcomes, followed by pairwise McNemar tests with Holm correction for multiple comparisons [181, 180, 182]. This complementary analysis was adopted in order to support pairwise conclusions using the matched correctness outcomes available for every sample.

For the operating-condition analysis, only the final selected pipeline was considered. Since the selected model produced one prediction per window, whereas the experimental conditions were defined at the original trial level, the window-level outputs belonging to the same trial were first aggregated into a single trial-level decision. Specifically, the confidence vectors of all windows belonging to the same trial were averaged, and the final predicted class for that trial was assigned as the class with the highest mean confidence. This probability-aware aggregation strategy is consistent with approaches used in segmented time-series classification when several fragments correspond to a single higher-level sample [183]. After aggregation, trial-level correctness was analyzed as a function of sequence, speed, and their interaction using a binomial generalized linear model with likelihood-ratio tests [184]. Statistical significance was assessed at $p < 0.05$, and multiple pairwise comparisons were corrected using the Holm procedure whenever applicable.

3.7 Results

This section presents the experimental results obtained from the proposed tactile hardness discrimination framework. The analyses include the outcomes of feature selection using PCA and RFE, model training and performance comparison across multiple classifiers, and the influence of experimental conditions on sensor responses. Subsequently, we evaluate the system's classification performance throughout the indentation sequence and in real-time operation, demonstrating the robustness and generalization capability of the proposed approach.

3.7.1 Feature subset definition and model selection

3.7.1.1 Feature subsets derived from PCA and RFE

As described in Section 3.6.2, six statistical features were extracted from each of the eight channels of every windowed trial, yielding a 48-dimensional feature vector. The two feature selection strategies, PCA-based ranking and RFE, were applied to both segmentation datasets, i.e., $\mathcal{D}_p$ and $\mathcal{D}_{100}$. In all configurations, the top 10 ranked features were retained.

To compactly denote these subsets, we use the notation $\mathcal{F}_{X-Y}$, where $X \in \{p, 100\}$ indicates the dataset (p for $\mathcal{D}_p$, 100 for $\mathcal{D}_{100}$) and $Y \in \{\text{PCA}, \text{RFE-SVM}, \text{RFE-RF}\}$ denotes the feature selection method. Feature labels in Table 3.2 follow the notation Feature_j, where $\text{Feature} \in \{\text{Skewness}, \text{Kurtosis}, \text{StD}, \text{PtP}, \text{Maximum}, \text{Minimum}\}$ denotes the descriptor type and $j \in \{1, \ldots, 8\}$ denotes the sensor channel on which that descriptor was computed. For example, Kurtosis_6 denotes the kurtosis of channel

Table 3.2: Selected 10-dimensional feature subsets obtained from PCA and RFE for the percentage-based dataset $\mathcal{D}_p$ and fixed-step dataset $\mathcal{D}_{100}$.

Subset	Selected features
$\mathcal{F}_{p\text{-}PCA}$	Kurtosis_6, Kurtosis_4, Kurtosis_8, Kurtosis_2, Kurtosis_3, Kurtosis_7, Skewness_6, Kurtosis_1, Skewness_4, Kurtosis_5
$\mathcal{F}_{100\text{-}PCA}$	Kurtosis_6, Kurtosis_4, Kurtosis_8, Kurtosis_2, Kurtosis_3, Kurtosis_7, Skewness_6, Kurtosis_1, Skewness_4, Kurtosis_5
$\mathcal{F}_{p\text{-}RFE\text{-}SVM}$	Skewness_1, PtP_1, Minimum_1, Skewness_2, Skewness_3, Skewness_4, Skewness_5, Kurtosis_5, Skewness_7, Skewness_8
$\mathcal{F}_{100\text{-}RFE\text{-}SVM}$	PtP_1, Minimum_1, PtP_2, PtP_4, Minimum_4, Maximum_4, PtP_6, Maximum_6, PtP_7, Maximum_7
$\mathcal{F}_{p\text{-}RFE\text{-}RF}$	Maximum_6, Skewness_8, Maximum_2, Minimum_1, StD_7, Maximum_7, Skewness_2, StD_1, Maximum_4, StD_2
$\mathcal{F}_{100\text{-}RFE\text{-}RF}$	Maximum_6, Skewness_8, Maximum_2, Minimum_1, StD_7, Maximum_7, Skewness_2, StD_1, Maximum_4, StD_2

6, PtP_1 the peak-to-peak amplitude of channel 1, and StD_7 the standard deviation of channel 7. The resulting 10-dimensional subsets are summarized in Table 3.2.

The feature names reported in Table 3.2 therefore correspond directly to the formally defined descriptors introduced in Section 3.6.2.

Notably, the PCA-based analysis yields the same subset of features for both segmentation schemes, suggesting that the principal sources of variance in the tactile data are largely invariant to the chosen windowing strategy. In contrast, the RFE-based subsets emphasize different combinations of features across channels, reflecting the estimator-dependent interaction between features and decision boundaries.

3.7.1.2 Cross-validation performance and model selection

The three ML models, i.e., SVM, RF, and SLFNN, were trained and tuned using the above feature subsets for both datasets. Model performance was assessed via five-fold stratified cross-validation, jointly stratified over experimental conditions.

Tables 3.3 and 3.4 summarize the mean accuracies and standard deviations obtained on the percentage-based and fixed-step datasets, respectively. On $\mathcal{D}_p$, SVM and RF trained on RFE-selected features achieved accuracies of 91.62 ± 2.1% and 94.65 ± 2.8%, whereas their counterparts trained on PCA-ranked features performed substantially worse, with accuracies of 75.97 ± 2.6% (SVM) and 77.68 ± 2.7% (RF). In contrast, the SLFNN trained on the PCA-ranked subset $\mathcal{F}_{p\text{-}PCA}$ reached the highest accuracy of **96.07** ± **1.2**%.

A similar trend was observed for the fixed-step dataset $\mathcal{D}_{100}$ (Table 3.4). With RFE-based features, SVM achieved 95.56 ± 1.04% accuracy and RF 93.15 ± 1.4%, while their performance with PCA-ranked features dropped to 84.45 ± 2.2% (SVM) and 80.95 ± 2.3% (RF). Again, the SLFNN trained on the PCA-based subset $\mathcal{F}_{100\text{-}PCA}$ achieved the best performance, with a mean accuracy of **97.07** ± **1.8**%.

These results reveal a clear pattern: RFE-based feature selection is beneficial for RF and SVM, which rely on discarding redundant or weakly informative variables. In contrast, the SLFNN, thanks to its nonlinear activation and capacity to model complex feature interactions, consistently achieves the highest accuracy when trained on the PCA-ranked subsets. Given its superior and stable performance across both segmentation schemes, the SLFNN with PCA-ranked features was selected as the reference model for the subsequent online classification.

Table 3.3: Performance comparison of RFE and PCA feature selection over percentage-based dataset $\mathcal{D}_p$. Values are mean accuracy ± standard deviation over five-fold cross-validation.

Model	Feature selection	Accuracy (%)
SVM	RFE	91.62 ± 2.1
Random Forest	RFE	94.65 ± 2.8
SVM	PCA	75.97 ± 2.6
Random Forest	PCA	77.68 ± 2.7
SLFNN	PCA	**96.07** ± **1.2**

To complement the descriptive cross-validation comparison reported above, formal statistical analyses were carried out on the fixed-step cumulative windowing pipelines used for model selection.

The inferential analysis (Table 3.5) confirmed that the observed differences among the five fixed-step pipelines were not merely descriptive. At the fold level, the Friedman

Table 3.4: Performance comparison of RFE and PCA feature selection over fixed-step dataset $\mathcal{D}_{100}$. Values are mean accuracy ± standard deviation over five-fold cross-validation.

Model	Feature selection	Accuracy (%)
SVM	RFE	95.56 ± 1.04
Random Forest	RFE	93.15 ± 1.4
SVM	PCA	84.45 ± 2.2
Random Forest	PCA	80.95 ± 2.3
SLFNN	PCA	**97.07** ± **1.8**

Table 3.5: Summary of the statistical analyses performed for the five fixed-step pipelines in Activity 2.

Analysis	**Test**	**Comparison**	**p-value**	**Conclusion**
Fold-level pipeline comparison	Friedman test	RF_PCA, SLFNN_PCA, SVM_PCA, SVM_RFE_SVM, RF_RFE_RF	4.99×10^{-4}	Significant overall difference among the five fixed-step pipelines
Sample-level pipeline comparison	Cochran's Q	Same five pipelines on matched out-of-fold predictions	< 0.001	Significant overall difference among the five pipelines
Post-hoc comparison	McNemar + Holm	SLFNN_PCA vs SVM_RFE_SVM	1.62×10^{-2}	SLFNN_PCA significantly outperformed SVM_RFE_SVM

test revealed a significant overall pipeline effect ($p = 4.99 \times 10^{-4}$), indicating that the five candidate pipelines did not perform equivalently. Since fold-level post-hoc comparisons are conservative with only five outer folds, a complementary sample-level analysis was also carried out using the matched out-of-fold predictions produced by all pipelines on the same held-out samples. Cochran's Q test was again highly significant ($p < 0.001$), confirming a global difference among the candidate pipelines. Post-hoc pairwise McNemar tests with Holm correction further showed that the best-performing configuration, SLFNN_PCA, significantly outperformed the second-best pipeline, SVM_RFE_SVM ($p_{\text{Holm}} = 0.0162$). These results provide formal statistical support for selecting the SLFNN trained on PCA-ranked features as the reference Activity 2 model.

3.7.1.3 Optimal SLFNN hyperparameters

After cross-validation, the optimal hyperparameters of the SLFNN were determined by selecting the configuration yielding the highest mean accuracy among the five folds for each dataset. For the percentage-based dataset $\mathcal{D}_p$, the best-performing SLFNN employed a single hidden layer with 100 neurons, a ReLU activation function, and an L2 regularization coefficient (α) of 0.001. The same network architecture was used for the fixed-step dataset $\mathcal{D}_{100}$, where α was slightly reduced to 0.0001 to improve generalization.

These optimal configurations, together with the PCA-ranked feature subsets $\mathcal{F}_{p\text{-}PCA}$ and $\mathcal{F}_{100\text{-}PCA}$, were adopted in the subsequent real-time testing phase, as they consistently provided the best balance between accuracy, stability, and computational efficiency across both segmentation schemes.

Taken together, these results identify the SLFNN trained on PCA-ranked features as the most suitable model for subsequent analyses, offering a good compromise between accuracy, stability, and implementation simplicity. In the following, we first examine how the experimental factors, indentation speed and target load, shape the underlying feature space, and then quantify how hardness discrimination performance evolves as a function of the observed portion of the tactile signal.

3.7.2 Effect of Experimental Conditions

The performance of the piezoelectric sensor array is strongly modulated by the indentation speed (IS) and target load (TL), which together shape the dynamic interaction between the fingertip and the object. Figure 3.3 illustrates representative sensor outputs under different combinations of IS and TL. In each subplot, the upper trace reports the force profile, while the lower traces show the voltage responses of the eight PVDF channels. At higher indentation speeds, the sensor signals exhibit more pronounced and faster fluctuations, consistent with the behavior of the skin-like interface and the rapid onset and release of contact. Conversely, variations in TL primarily affect the overall response amplitude: larger loads lead to stronger charge generation in the piezoelectric material and hence higher voltage excursions. These examples confirm that both speed and load systematically alter the temporal and amplitude characteristics of the tactile signals, justifying their controlled variation during data collection.(Section 3.4)

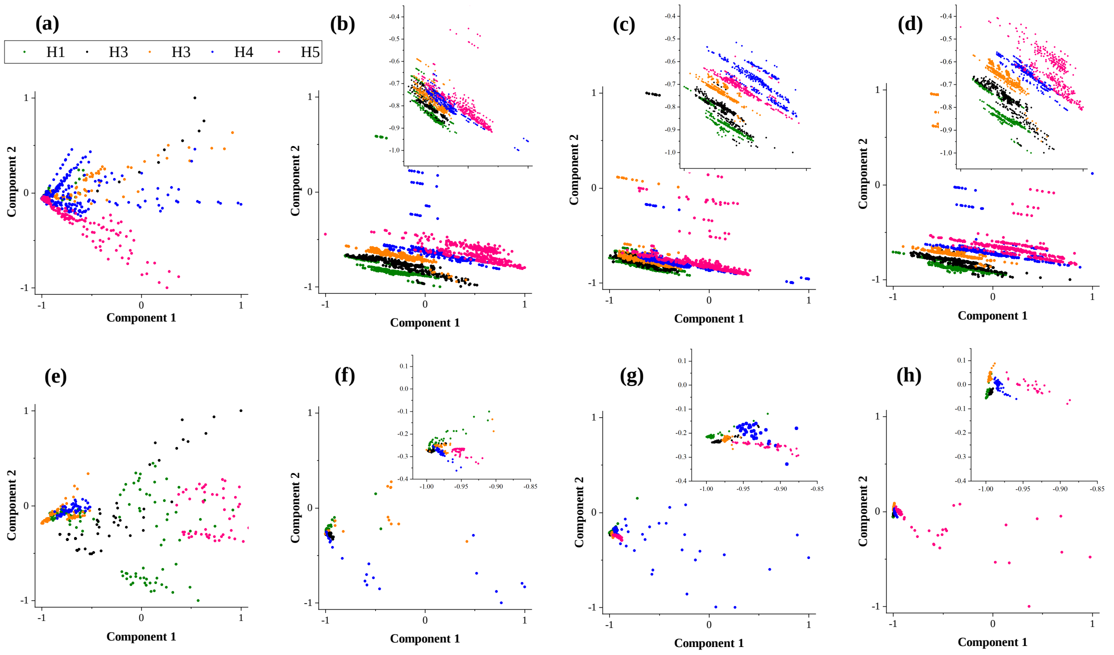

Figure 3.8: This figure represents the PCA biplot of the data representation over all the conditions and presented as in subplot (a) PCA over the whole data in (b) Indentation speed 1 (c) Indentation speed 2 (d) Indentation speed 3 (e) Target load 1 (f) Target load 2 (g) Target load 3 (h) Target load 4

To quantify how these factors influence hardness separability in the feature space, a PCA was performed on the Kurtosis_6, and Kurtosis_4 of $\mathcal{F}_{p\text{-}PCA}$. This data was illustrated in Fig. 3.8. Subplot (a) considers all windows jointly and shows that the five hardness classes form partially overlapping yet distinguishable clusters in the first two principal components. The data was organized along an elongated manifold, with hardness levels distributed along this structure rather than being linearly separable, highlighting the need for nonlinear classifiers such as the SLFNN.

Subplots (b)–(d) isolate trials acquired at fixed indentation speeds, while preserving the color-coding by hardness. At the lowest speed (Fig. 3.8(b)), hardness clusters are compact but partially collapsed onto each other, indicating limited separability. As the speed increases (Fig. 3.8(c)–(d)), the clusters progressively fan out along the principal components, and the intra-class spread becomes more structured. This suggests that higher indentation speeds enhance the discriminative content of the features, likely because faster interactions emphasize dynamic differences in the mechanical response of objects with different hardness.

A similar trend is observed when the analysis is conditioned on target load (Fig. 3.8(e)–(h)). For lower TL values, the clusters corresponding to different hardness levels remain close and strongly overlapping, whereas at higher loads, the class-specific clusters become more compact and better separated in the PCA plane. In particular, the highest loads (Fig. 3.8(g)–(h)) yield well-structured clusters with limited overlap, indicating that stronger finger–object interactions amplify hardness-dependent variations captured by the piezoelectric sensors.

Overall, the PCA visualizations confirm that both IS and TL modulate the geometry of the feature space: increasing the parameter tends to improve the separation between hardness classes. This supports the experimental design, where multiple force and speed conditions were deliberately explored to enrich the dataset and enable robust hardness discrimination across a range of interaction scenarios.

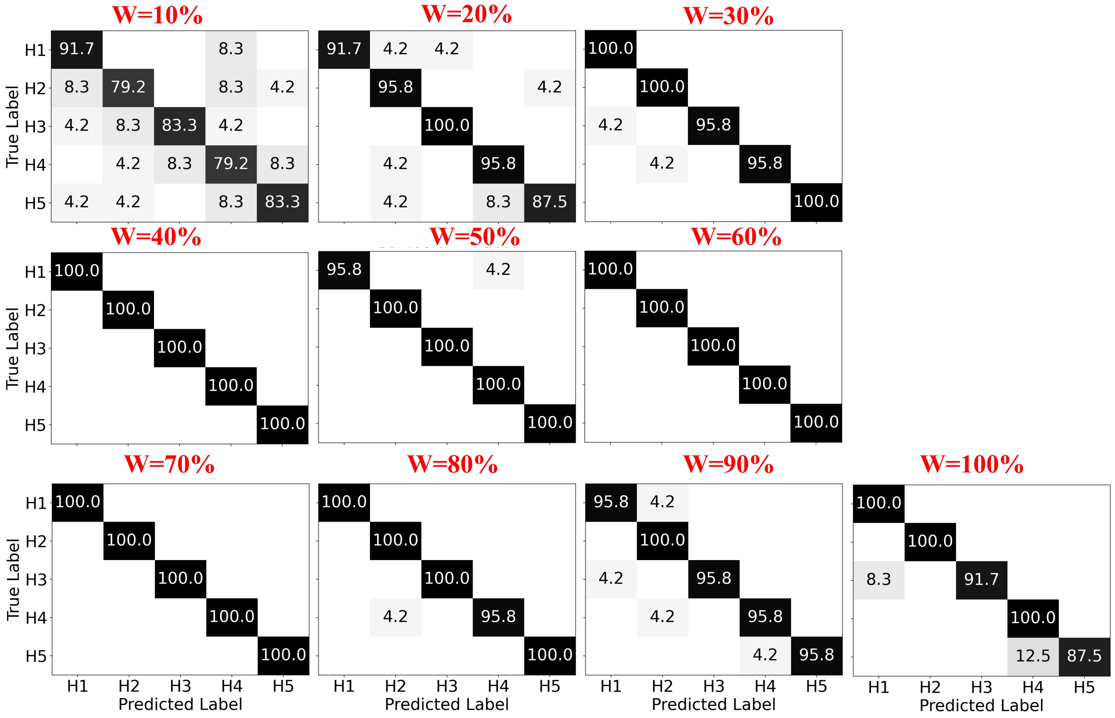

W=10%	H1	H2	H3	H4	H5
H1	91.7			8.3	
H2	8.3	79.2		8.3	4.2
H3	4.2	8.3	83.3	4.2	
H4		4.2	8.3	79.2	8.3
H5	4.2	4.2		8.3	83.3

W=20%	H1	H2	H3	H4	H5
H1	91.7	4.2	4.2		
H2		95.8			4.2
H3			100.0		
H4		4.2		95.8	
H5		4.2		8.3	87.5

W=30%	H1	H2	H3	H4	H5
H1	100.0				
H2		100.0			
H3	4.2		95.8		
H4		4.2		95.8	
H5					100.0

W=40%	H1	H2	H3	H4	H5
H1	100.0				
H2		100.0			
H3			100.0		
H4				100.0	
H5					100.0

W=50%	H1	H2	H3	H4	H5
H1	95.8			4.2	
H2		100.0			
H3			100.0		
H4				100.0	
H5					100.0

W=60%	H1	H2	H3	H4	H5
H1	100.0				
H2		100.0			
H3			100.0		
H4				100.0	
H5					100.0

W=70%	H1	H2	H3	H4	H5
H1	100.0				
H2		100.0			
H3			100.0		
H4				100.0	
H5					100.0

W=80%	H1	H2	H3	H4	H5
H1	100.0				
H2		100.0			
H3			100.0		
H4		4.2		95.8	
H5					100.0

W=90%	H1	H2	H3	H4	H5
H1	95.8	4.2			
H2		100.0			
H3	4.2		95.8		
H4		4.2		95.8	
H5				4.2	95.8

W=100%	H1	H2	H3	H4	H5
H1	100.0				
H2		100.0			
H3	8.3		91.7		
H4				100.0	
H5				12.5	87.5

Figure 3.9: Confusion Matrices for model behavior and accuracy over each window of $\mathcal{D}_p$

3.7.3 Effect of percentage-based segmentation

To investigate how hardness information accumulates over the course of the indentation, we first analyzed classification performance on the percentage-based dataset $\mathcal{D}_p$, where each window contains a fixed fraction of the filtered signal (10%, 20%, . . . , 100%). For each percentage, a confusion matrix was computed over the test set, expressing for every true hardness level (H1–H5) the percentage of windows assigned to each predicted class. The resulting matrices are summarized in Fig. 3.9 for window percentages from 10% to 100%.

At 10% of the signal, the model already achieves reasonably high per-class accuracies, with diagonal entries of 91.7% (H1), 79.2% (H2), 83.3% (H3), 79.2% (H4), and 83.3% (H5), corresponding to a mean class accuracy of approximately 83.3%. Most errors occur between adjacent or near-adjacent hardness levels. For instance, H2 is sometimes misclassified as H1, H4, or H5, while H3 and H4 are occasionally confused with their neighbouring classes. This pattern reflects the fact that, in the very early phase of contact, only limited dynamic information is available, and the tactile signatures of different hardness levels are still partially overlapping.

When 20% of the signal is considered, the class-specific accuracies increase markedly to 91.7% (H1), 95.8% (H2), 100.0% (H3), 95.8% (H4), and 87.5% (H5), with a mean of about 94.2%. At this stage, the confusion for H3 disappears completely, and H2 and H4 are correctly classified in the vast majority of cases. Residual errors mainly involve H1 being occasionally mistaken for H2 or H3, and H5 being confused with H2 or H4, again consistent with local ambiguities between intermediate hardness levels.

From 30% onward, hardness discrimination becomes almost perfect. At 30%, the diagonal entries reach 100.0% (H1), 100.0% (H2), 95.8% (H3), 95.8% (H4), and 100.0% (H5), yielding a mean class accuracy of 98.3%. Only two small confusion terms remain: a 4.2% misclassification of H3 as H1 and a 4.2% misclassification of H4 as H2. At 40% of the signal, the confusion matrix becomes perfectly diagonal, with 100% correct classification for all five hardness levels. This behavior persists at 60% and 70%, where all classes are again classified with 100% accuracy, indicating that once a sufficient portion of the indentation has been observed, the feature representation extracted from $\mathcal{D}_p$ is fully discriminative for hardness.

For intermediate percentages between 40% and 100%, the matrices remain very close to the identity. At 50%, a small fraction (4.2%) of H1 windows is misclassified as H4, leading to a mean class accuracy of 99.2%. At 80%, a single off-diagonal term appears, with 4.2% of H4 windows assigned to H2 (mean accuracy 99.2%). At 90%, minor confusions re-emerge: 4.2% of H1 windows are predicted as H2, 4.2% of H3 as H1, 4.2% of H4 as H2, and 4.2% of H5 as H4, resulting in a mean class accuracy of 96.6%. Finally, at 100% of the signal, the diagonal remains high (100.0% for H1, H2, and H4; 91.7% for H3; and 87.5% for H5), but a small number of H3 windows are assigned to H1 (8.3%) and some H5 windows to H4 (12.5%), yielding a mean class accuracy of 95.8%.

Two main conclusions emerge from this analysis. First, hardness discrimination improves rapidly with the percentage of the signal: moving from 10% to 30% of the indentation increases the mean class accuracy from approximately 83% to over 98%, and from 40% to 70% the classifier achieves perfect separation of the five hardness levels on $\mathcal{D}_p$. This indicates that the most informative portion of the tactile interaction is concentrated in the early-to-middle phase of the indentation, and that only a relatively small fraction of the signal is needed to reach near-optimal performance. Second, the small fluctuations observed at higher percentages (80–100%) suggest that the very late part of the signal, which may include release dynamics and variability at the end of the interaction, can introduce slight ambiguity for some classes (particularly H3 and H5). Nonetheless, these effects remain limited, and overall hardness classification remains highly accurate across all percentages once at least 20–30% of the indentation has been observed.

3.7.4 Effect of fixed-step incremental segmentation

We next analyzed the behavior of the classifier on the fixed-step dataset $\mathcal{D}_{100}$, where each window aggregates an increasing number of samples in steps of 100, starting from contact onset. It is important to note that the notion of "percentage" in $\mathcal{D}_{100}$ differs from that used for the percentage-based dataset $\mathcal{D}_p$. In $\mathcal{D}_p$, a window at p% contains exactly p% of the samples of the filtered signal, so the window length is explicitly defined as a fraction of the trial duration. In contrast, $\mathcal{D}_{100}$ is constructed by cumulatively aggregating fixed-size segments of 100 samples: the first window contains the first 100

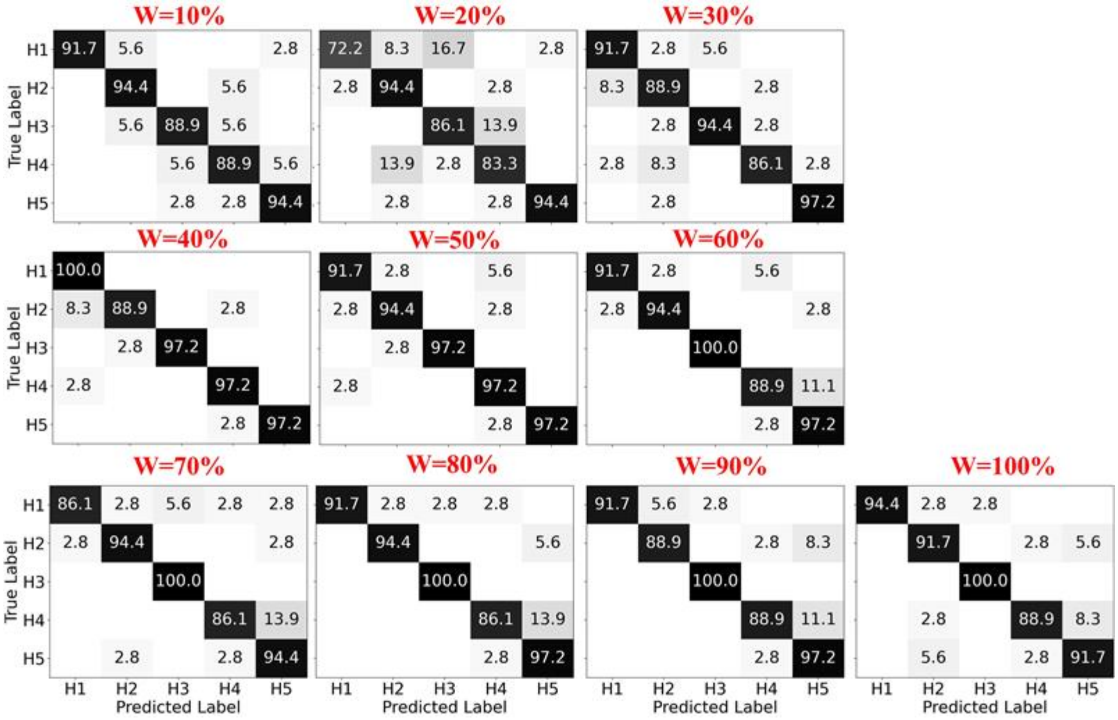

W=10%

	H1	H2	H3	H4	H5
H1	91.7	5.6			2.8
H2		94.4		5.6	
H3		5.6	88.9	5.6	
H4			5.6	88.9	5.6
H5			2.8	2.8	94.4

W=20%

	H1	H2	H3	H4	H5
H1	72.2	8.3	16.7		2.8
H2	2.8	94.4		2.8	
H3			86.1	13.9	
H4		13.9	2.8	83.3	
H5		2.8		2.8	94.4

W=30%

	H1	H2	H3	H4	H5
H1	91.7	2.8	5.6		
H2	8.3	88.9		2.8	
H3		2.8	94.4	2.8	
H4	2.8	8.3		86.1	2.8
H5		2.8			97.2

W=40%

	H1	H2	H3	H4	H5
H1	100.0				
H2	8.3	88.9		2.8	
H3		2.8	97.2		
H4	2.8			97.2	
H5				2.8	97.2

W=50%

	H1	H2	H3	H4	H5
H1	91.7	2.8		5.6	
H2	2.8	94.4		2.8	
H3		2.8	97.2		
H4	2.8			97.2	
H5				2.8	97.2

W=60%

	H1	H2	H3	H4	H5
H1	91.7	2.8		5.6	
H2	2.8	94.4			2.8
H3			100.0		
H4				88.9	11.1
H5				2.8	97.2

W=70%

	H1	H2	H3	H4	H5
H1	86.1	2.8	5.6	2.8	2.8
H2	2.8	94.4			2.8
H3			100.0		
H4				86.1	13.9
H5		2.8		2.8	94.4

W=80%

	H1	H2	H3	H4	H5
H1	91.7	2.8	2.8	2.8	
H2		94.4			5.6
H3			100.0		
H4				86.1	13.9
H5				2.8	97.2

W=90%

	H1	H2	H3	H4	H5
H1	91.7	5.6	2.8		
H2		88.9		2.8	8.3
H3			100.0		
H4				88.9	11.1
H5				2.8	97.2

W=100%

	H1	H2	H3	H4	H5
H1	94.4	2.8	2.8		
H2		91.7		2.8	5.6
H3			100.0		
H4		2.8		88.9	8.3
H5		5.6		2.8	91.7

Figure 3.10: Confusion Matrices for model behavior and accuracy over each window of $\mathcal{D}_{100}$

samples, the second window the first 200 samples, and so on, up to a trial-specific last window index K_i. For the analysis in Fig. 3.10, the percentages (10%, 20%, . . . , 100%) are computed per trial as a fraction of K_i, and the corresponding cumulative window is selected (e.g., the 10% window is approximately $\lfloor 0.1\, K_i \rfloor$). As a consequence, a 10% window in $\mathcal{D}_{100}$ always corresponds to an early portion of the interaction, and a 100% window to the entire interaction, but the exact number of samples per selected window can vary across trials and scenarios (due to differences in indentation speed and target load). The resulting confusion matrices was computed over the test set (Fig. 3.10). Each matrix reports, for a given true hardness level (H1–H5), the percentage of windows assigned to each predicted class.

At 10% of the available windows, the classifier already achieves high per-class accuracies: 91.7% (H1), 94.4% (H2), 88.9% (H3), 88.9% (H4), and 94.4% (H5), corresponding to a mean class accuracy of 91.7%. The residual errors are local and mostly involve neighboring hardness levels. For example, 5.6% of H1 windows are misclassified as H2, 5.6% of H2 as H4, and 5.6% of H3 as either H2 or H4, while H4 is occasionally confused with H3 or H5 (both at 5.6%). These near-miss errors indicate that, even at the very beginning of the incremental process, the fixed-size windows already contain enough information to discriminate hardness reasonably well, although some ambiguity remains between adjacent stiffness levels.

At 20%, the mean class accuracy drops to 86.1%, mainly due to a degradation in H1 classification (72.2% correct). In this regime, 16.7% of H1 windows are assigned to H3 and 8.3% to H2, while 13.9% of H3 windows are misclassified as H4 and 13.9% of H4 windows as H2. These confusions suggest that, for some trials, the intermediate portion of the indentation does not yet fully separate the elastic responses of soft and medium-hard objects, and that transient dynamics can temporarily blur the distinction between neighboring hardness levels.

From 30% onward, the overall performance stabilizes at a high level. At 30%, the per-class accuracies are 91.7% (H1), 88.9% (H2), 94.4% (H3), 86.1% (H4), and 97.2% (H5), with a mean accuracy of 91.7%. The confusion structure remains largely local, with small fractions of H1 predicted as H2 or H3, H2 occasionally mapped to H4, and H4 rarely misclassified as H1 or H2.

The best performance is achieved around 40–50% of the available windows. At 40%, the diagonal entries reach 100.0% (H1), 88.9% (H2), 97.2% (H3), 97.2% (H4), and 97.2% (H5), yielding a mean class accuracy of 96.1%. Misclassifications are limited to a small fraction of H2 windows assigned to H1 or H4, and isolated 2.8% confusions between neighboring classes. At 50%, the mean accuracy remains comparably high (95.5%), with 91.7% (H1), 94.4% (H2), and 97.2% for H3–H5. Here, the most frequent residual errors involve H1 being classified as H4 (5.6%) and symmetric low-rate confusions between H1 and H4, as well as between H3 and H2 or H5 and H4 (all at 2.8%).

For larger percentages (60–100%), the performance remains consistently high, with mean class accuracies between 92.2% and 94.4%. Class H3 becomes almost perfectly recognized, reaching 100% accuracy from 60% onward, indicating that the corresponding tactile signatures are particularly robust once a sufficient portion of the incremental signal has been accumulated. In contrast, most of the residual confusion concentrates on the harder classes H4 and H5. At 70–90%, 11.1–13.9% of H4 windows are predicted as H5, and 2.8–8.3% of H5 windows as H2 or H4. At 100%, the diagonal entries are 94.4%

(H1), 91.7% (H2), 100.0% (H3), 88.9% (H4), and 91.7% (H5), with a mean accuracy of 93.3%. The remaining errors are again localized and primarily involve confusion between the hardest levels (H4, H5) and, to a lesser extent, between H2 and H5.

Overall, the fixed-step incremental windowing scheme yields consistently high hardness classification accuracy across all percentages once the first few windows have been observed. Unlike the percentage-based segmentation, where performance sharply transitions to near-perfect accuracy between 20% and 40%, the fixed-step scheme exhibits a more gradual evolution: the model is already strong at 10%, shows a transient dip at 20%, and then stabilizes above 90% mean accuracy from 30% onward. The residual errors are dominated by near-miss confusions between adjacent hardness levels, especially among the hardest classes, indicating that the SLFNN reliably captures the global structure of the incremental tactile response while remaining sensitive to subtle similarities between neighboring stiffness values.

3.7.5 Comparison between percentage-based and fixed-step segmentation

Although the percentage indices reported for $\mathcal{D}_p$ and $\mathcal{D}_{100}$ are defined differently, they can be meaningfully compared in terms of the relative progression of the indentation. In $\mathcal{D}_p$, the window at p% contains exactly p% of the samples of the filtered signal, so the time span covered by the window is explicitly proportional to the trial duration. In $\mathcal{D}_{100}$, the windows are constructed cumulatively in fixed increments of 100 samples, and the indices used to build the confusion matrices are obtained by selecting windows at fixed fractions (10%, 20%, ..., 100%) of the last cumulative window index of each trial. Thus, while the number of samples in the selected window is not identical across all trials in $\mathcal{D}_{100}$, the percentage still reflects an early, intermediate, or late stage of the interaction for each trial.

From this perspective, the two segmentation schemes provide complementary views on how hardness information accumulates over the course of the indentation. In the percentage-based dataset $\mathcal{D}_p$, classification performance improves very rapidly: the mean class accuracy increases from about 83% at 10% to over 98% at 30%, and becomes essentially perfect (100% mean accuracy) between 40% and 70%. This suggests that a relatively small fraction of the signal, expressed as a percentage of the trial duration, is sufficient to fully separate the hardness classes when features are computed on fixed relative portions of the indentation.

In the fixed-step dataset $\mathcal{D}_{100}$, the evolution of performance is more gradual but also more robust across percentages. The classifier already achieves high accuracy at 10% (about 92%), experiences a transient decrease at 20%, and then stabilizes above 90% mean accuracy from 30% onward, with a maximum around 40–50%. Residual errors in $\mathcal{D}_{100}$ are consistently dominated by near-miss confusions between neighboring hardness levels, especially among the hardest classes (H4 and H5), indicating that the incremental 100-sample windows capture the global structure of the tactile response while being slightly more sensitive to local variations in contact timing and release.

Taken together, these results show that, despite the different definitions of "percentage", both segmentation schemes reveal a coherent picture: hardness-discriminative information emerges early in the indentation and progressively stabilizes as more of the signal is accumulated. The percentage-based windows emphasize performance as a function of the relative fraction of the signal, whereas the fixed-step incremental windows emphasize performance as a function of discrete cumulative observations. In both cases, reliable hardness classification is achieved well before the end of the interaction, which is a desirable property for real-time deployment.

3.7.6 Robustness of the selected model across Target Loads and Indentation Speed conditions

To further assess the robustness of the selected Activity 2 model under the operating conditions used during data collection, the held-out test predictions were analyzed as a function of sequence and indentation speed.

For this analysis, the window-level outputs of the selected SLFNN_PCA model were aggregated into trial-level decisions by averaging the confidence vectors of all windows belonging to the same original trial and assigning the final label as the class with the highest mean confidence. This resulted in 11,869 window-level predictions being summarized into 180 trial-level decisions, with an overall trial-level accuracy of 97.78%. Descriptively, the trial-level accuracies remained very high across all sequence and speed levels. Sequence levels 1 and 2 achieved 95.56%, whereas sequence levels 3

Table 3.6: Trial-level performance of the selected Activity 2 model (SLFNN_PCA) across sequence and speed conditions after aggregation of window-level predictions by mean confidence.

Factor	**Level**	**Accuracy (%)**
Sequence	1	95.56
Sequence	2	95.56
Sequence	3	100.00
Sequence	4	100.00
Speed	150	98.33
Speed	300	96.67
Speed	450	98.33

and 4 reached 100%. For speed, the trial-level accuracies were 98.33% at 150, 96.67% at 300, and 98.33% at 450.

To determine whether these descriptive variations corresponded to statistically meaningful condition effects, a binomial generalized linear model with likelihood-ratio tests was applied. The analysis did not reveal a statistically significant main effect of target loads ($p = 0.1307$), indentation speeds ($p = 0.7856$), or their interaction ($p = 0.8285$) on trial-level correctness. Therefore, although small descriptive differences can be observed across operating conditions, no statistically significant evidence of condition-dependent degradation was found. These results support the robustness of the selected SLFNN_PCA pipeline under the tested sequence and speed settings.

3.7.7 Analysis of probabilistic outputs

To complement the window-level accuracy and confusion matrices reported in the previous sections, we further analyzed the probability vectors produced by the SLFNN classifier on $\mathcal{D}_{100}$. Rather than considering only the most probable class, this analysis exploits the full distribution over the five hardness levels (H1–H5). It focuses on four aspects: (i) Top-k accuracy and “near-miss” behavior, (ii) calibration of the predicted probabilities, (iii) distribution of confidence across hardness classes, and (iv) rank of the true class among the model's probabilistic outputs for misclassified windows. All results in this section are computed over 11,869 windows in the test set.

3.7.7.1 Top-k accuracy and near-miss behavior

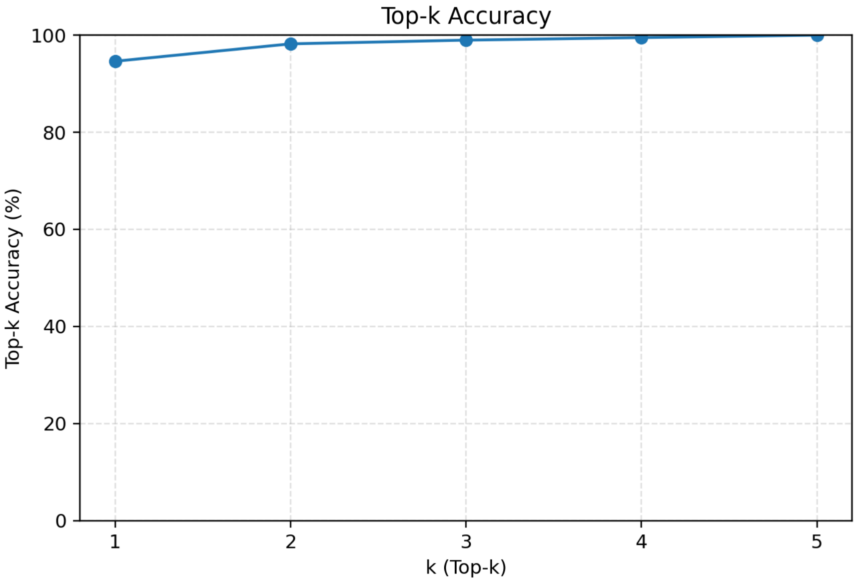

Figure 3.11: Top-k accuracy of the SLFNN classifier on the fixed-size window dataset.

Figure 3.11 reports the Top-k accuracy of the classifier computed over all test windows. For each window, the model outputs a probability distribution over the five hardness classes. These probabilities are sorted in descending order, and a prediction is counted as correct under the Top-k criterion if the ground-truth class appears among the k most probable outputs.

The model achieves a Top-1 accuracy of 94.6%, corresponding to the standard notion of accuracy when only the most probable class is considered. When Top-2 is used, the accuracy increases to 98.2%, and further to 99.0% for Top-3. These results indicate that, even when the Top-1 prediction is wrong, the correct hardness level often appears among the most probable alternatives.

To better characterize these "near-miss" cases, we examined only the misclassified windows, i.e., those failing the Top-1 criterion ($N = 638$). For each such window, we computed the rank of the true class within the sorted probability vector (rank 1 corresponds to the chosen class, rank 2 to the second-highest probability, and so on). The distribution of these ranks is shown in Fig. 3.14. Out of the 638 misclassified windows, the true class appears as the second most probable class in 425 cases, corresponding to 66.6% of all misclassifications. The true class is ranked third in 91 cases (14.3%), fourth in 62 cases (9.7%), and fifth in 60 cases (9.4%).

This distribution confirms that the majority of errors correspond to near-miss events: in roughly two-thirds of misclassified windows the classifier assigns the highest probability to an incorrect hardness level but still places the true class immediately second in its preference ordering. Only a small fraction of errors ($< 20\%$ in total) involve the true class being relegated to the bottom of the ranking (ranks 4 or 5). From an application perspective, this behavior is desirable, as it suggests that a Top-2 decision rule, a confidence-based rejection strategy, or a downstream fusion mechanism could recover a large portion of nominal Top-1 errors by exploiting the information contained in the full probability vector.

3.7.7.2 Calibration of predicted probabilities

While accuracy quantifies how often the predicted class is correct, probability calibration assesses how well the predicted confidence values reflect the true likelihood of correctness. To evaluate calibration, we constructed a reliability diagram (Fig. 3.12). For each window, we considered the maximum predicted probability (i.e., the confidence assigned to the chosen class) and grouped windows into confidence bins (0.3–0.4, 0.4–0.5, . . . , 0.9–1.0). Within each bin, we computed the mean predicted confidence and the empirical accuracy (fraction of correct predictions).

The resulting curve lies close to the identity line, indicating that the model is reasonably well calibrated. For low-confidence predictions (0.3–0.4), the mean confidence is 36.8% while the observed accuracy is 26.3% ($N = 19$ windows), reflecting substantial uncertainty as expected. As confidence increases, accuracy also improves: in the 0.8–0.9 bin, the model reports an average confidence of 85.6%, with an empirical accuracy of 77.5% ($N = 298$). The majority of windows (92.6%) fall into the highest confidence bin (0.9–1.0), where the mean confidence is 99.6%, and the accuracy reaches 97.1% ($N = 10\,995$).

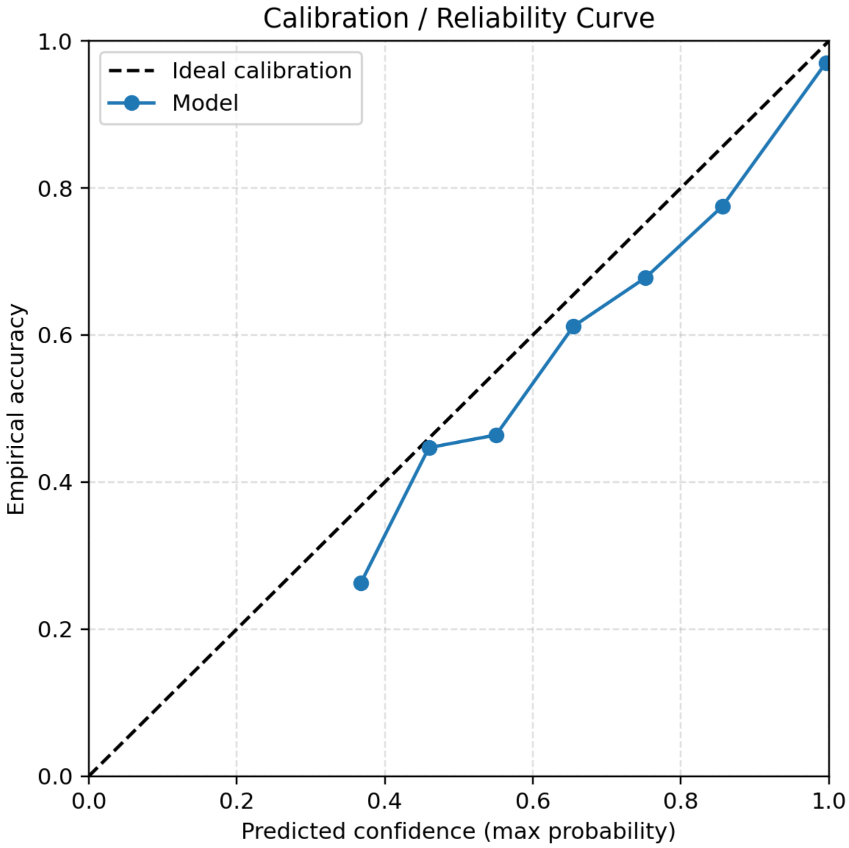

Figure 3.12: Reliability diagram of the SLFNN classifier.

Overall, the model exhibits mild overconfidence with an expected calibration error of approximately 2.9 percentage points; that is, on average the predicted confidence slightly exceeds the true accuracy. Nevertheless, the monotonic relationship between confidence and accuracy demonstrates that the predicted probabilities are informative: higher confidence reliably indicates a higher probability of correct classification. In an online deployment, these properties could be exploited to define confidence thresholds for triggering “uncertain” decisions, requesting additional information, or adapting the control strategy of the prosthetic device.

3.7.7.3 Per-class confidence profiles

Figure 3.13: Mean confidence assigned to the true hardness class for each hardness level (H1–H5), with error bars representing the standard deviation across windows.

To investigate how the model's confidence depends on the hardness level, we analyzed the probability assigned to the true class for each window and aggregated these values per hardness class (H1–H5). For every window, we extracted the probability associated with the ground-truth class, expressed it as a percentage, and computed summary statistics across all windows belonging to the same hardness level. Figure 3.13 shows the mean confidence and standard deviation for each class.

The model assigns consistently high confidence to the true class across all hardness levels, with mean values of **95.2%** for H1, **91.3%** for H2, **96.2%** for H3, **87.8%** for H4, and **97.6%** for H5. The highest confidence is observed for the extreme hardness levels (H1 and H5), whereas intermediate levels (H2 and H4) exhibit lower mean confidence and larger variability (standard deviations of approximately 25–29 percentage points).

These confidence profiles are consistent with the per-class accuracies derived from the confusion matrix: 96.1% (H1), 92.2% (H2), 97.1% (H3), 89.0% (H4), and 98.2%

(H5). Thus, the classifier is both more accurate and more confident for very soft and very hard objects, while intermediate hardness levels are intrinsically more ambiguous for the model. From a tactile perception standpoint, this behavior is plausible: mid-range stiffness values are closer to each other in terms of mechanical response and therefore more prone to overlap in the sensor space. In line with this, when separating windows by correctness, the mean probability assigned to the true class is approximately 98.4% for correctly classified windows and only 12.2% for misclassified ones, confirming that low confidence on the true class strongly correlates with errors.

3.7.7.4 Rank of the true class among misclassified predictions

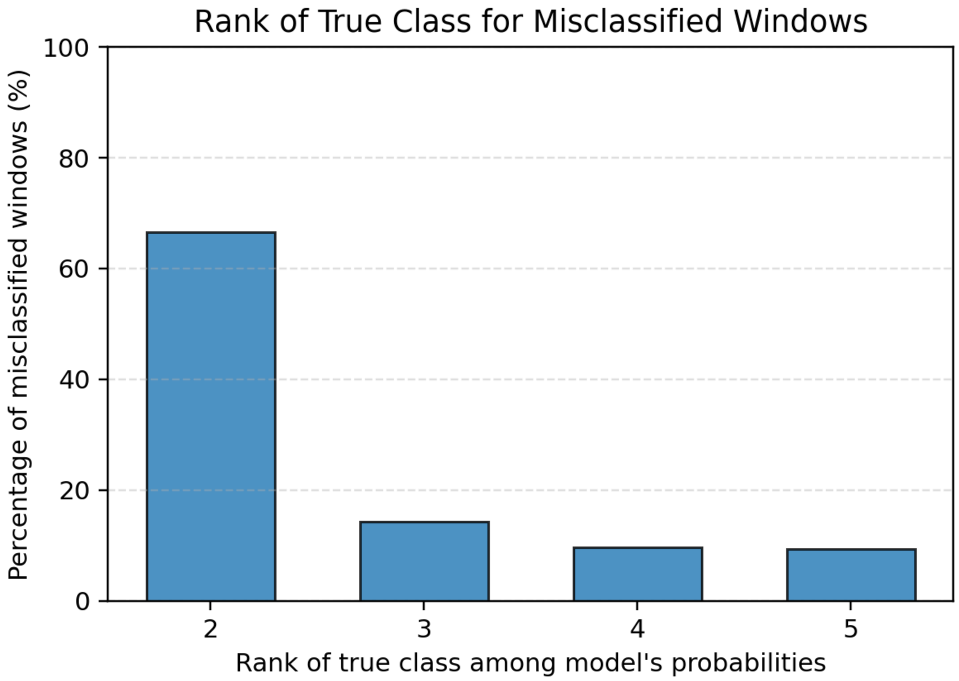

Figure 3.14: Distribution of the rank of the true hardness class among the model's predicted probabilities for misclassified windows.

Finally, we examined in more detail the position of the true class within the model's probability ranking for misclassified windows. For each misclassified window, we sorted the class probabilities in descending order and recorded the rank of the true hardness class. The corresponding distribution is reported in Fig. 3.14 and numerically summarized in Sec. 3.7.7.1.

Out of 638 misclassified windows, the true class appears as the second most probable class in 425 cases (66.6%), as the third most probable class in 91 cases (14.3%), and as the fourth and fifth most probable classes in 62 (9.7%) and 60 (9.4%) cases, respectively. Only a small fraction of errors therefore, correspond to situations where the model assigns very low probability to the true hardness level. This behaviour supports the view that most misclassifications are localized confusions between neighbouring hardness levels, rather than random or inconsistent errors.

3.7.7.5 Summary of probabilistic behaviour

Taken together, these probabilistic analyses show that the SLFNN provides not only high Top-1 accuracy, but also meaningful and interpretable confidence estimates. The Top-k accuracy and rank distributions reveal that most errors are near misses in which the true hardness level remains among the highest-ranked alternatives. The calibration analysis indicates that the predicted probabilities are reasonably well aligned with empirical accuracy, especially in the high-confidence regime that dominates the online operation, albeit with a slight tendency toward overconfidence. Finally, the per-class confidence profiles highlight that very soft and very hard objects are recognized with higher confidence and accuracy than intermediate stiffness levels, which are intrinsically more ambiguous from a tactile sensing perspective. These properties suggest that the probabilistic outputs of the classifier can be reliably exploited to implement confidence-aware decision strategies and to quantify uncertainty in real-time hardness discrimination.

3.8 Discussion

This activity investigated whether high-accuracy hardness recognition can be achieved in a way that is consistent with incremental tactile perception, i.e., by making reliable decisions before the indentation is fully completed, and under interaction variability induced by target load (TL) and indentation speed (IS). The results support three main technical insights. First, although 48 per-window features were extracted (six statistics per channel across eight sensors), the model-selection results show that the best performance was obtained with a compact 10-feature subset and a lightweight classifier, i.e., an SLFNN trained on PCA-ranked features, reaching high cross-validated accuracy on both segmentation schemes. Moreover, that ranking was further supported by formal statistical analysis on the fixed-step pipelines, which confirmed a significant overall difference among the candidate configurations and identified the SLFNN_PCA pipeline as the best-performing solution. This is practically important because compact feature sets reduce latency and improve stability in online pipelines, while still capturing interaction-relevant dynamics (here, mainly higher-order statistics such as kurtosis and skewness across specific channels). The behavior of PCA vs. RFE across models is also consistent with what is expected from the learning bias of the estimators: wrapper methods such as RFE are often advantageous for margin-based or tree-based models because they directly optimize discrimination under the estimator's inductive assumptions, whereas the SLFNN can exploit nonlinear interactions among the PCA-ranked variables even when the ranking itself is variance-driven rather than label-driven [174, 168]. Second, the segmentation analysis directly answers the "how early can we decide?" question: with percentage-based windows, accuracy rises sharply and becomes essentially perfect over a mid-portion of the signal (typically after observing only a fraction of the interaction), while the fixed-step incremental windows produce a more gradual but consistently strong performance curve and, crucially, correspond to a segmentation mechanism that is implementable online because it does not require knowing the final signal length in advance. Together, these two views indicate that hardness-discriminative information concentrates early-to-mid indentation, and that incremental accumulation in discrete steps (here, 100-sample cumulative windows) is a robust strategy for real-time operation. This "early decision" capability is a central novelty relative to a large body of tactile hardness work that evaluates classification on complete interactions and/or under more constrained motion policies, including robotic-gripper studies that focus on global accuracy or resource trade-offs without explicitly characterizing decision timing as a function of observed signal fraction [166, 161, 185, 95, 28]. Third, beyond Top-1 accuracy, the probabilistic-output analysis adds an additional layer of deployability: most errors are *near misses* (true class often ranked second), confidence monotonically tracks correctness, and calibration is reasonably good with only mild overconfidence. This matters because calibrated (or calibratable) posteriors enable confidence-aware controllers (e.g., delaying a control action until confidence exceeds a threshold, or fusing Top-2 hypotheses with other sensory cues), which is increasingly recognized as a practical requirement when learned tactile perception is integrated into real robot or prosthesis pipelines [186, 28]. Compared with prior work on the same TPU object family implemented on embedded platforms [166], our emphasis is not hardware benchmarking (since this chapter's pipeline runs on a PC), but rather the algorithmic characterization that is directly actionable for future embedded deployment: (i) incremental windowing that supports early decisions, (ii) compact features computable online without storing full histories, and (iii) interpretable probabilistic outputs that expose uncertainty. At the same time, the present study remains limited to controlled indentation trajectories; extending these results to less constrained interactions (variable contact geometry, incidental slip, non repeatable trajectories) requires validation under more ecological conditions. This limitation directly motivates the next step of the thesis (Chapter 4), where hardness discrimination is moved from the controlled Cartesian robot to a less controlled prosthetic-hand setting, so that the algorithmic insights established here (robustness to TL/IS variability, early-decision behavior, and confidence structure) can be tested under realistic grasp dynamics.

3.9 Conclusions

This chapter presented a PVDF-based biomimetic tactile sensing system integrated into a silicone finger cap and combined it with machine-learning pipelines to discriminate five discrete hardness levels of 3D-printed TPU objects under controlled, repeatable indentation while varying target load and indentation speed. In Activity 1, we established a low-complexity baseline by selecting the most active channel, extracting a small set of statistical descriptors, and showing that conventional classifiers (SVM/SLFNN/KNN) can achieve high accuracy even with very compact representations; in particular, strong performance was obtained using only a handful of features from a single sensor, and near-perfect performance was possible when operating conditions were fixed at favorable TL/IS points [173]. This first activity therefore quantified the attainable accuracy–complexity trade-off when the sensing problem is simplified to a single informative channel.

Activity 2 then moved from offline, trial-level evaluation to an incremental perception setting that better matches how tactile information becomes available in practice. By filtering and segmenting each interaction into progressively increasing windows, extracting features that can be computed online, and systematically comparing feature-selection strategies and lightweight classifiers, we demonstrated that reliable hardness recognition can be achieved *before* the indentation is completed. Across the two windowing strategies, hardness-discriminative information emerged early-to-mid interaction, and the fixed-step cumulative segmentation provided an online-feasible mechanism with consistently high accuracy. Finally, probabilistic-output analyses (Top-k, rank distribution, and calibration) showed that the classifier's confidence is informative and that most errors are near-miss confusions between neighboring hardness levels, enabling future confidence-aware decision policies [186]. Overall, the chapter contributes both (i) empirical evidence that PVDF-based tactile sensing can remain accurate under TL/IS variability, and (ii) an incremental classification framework that exposes when and how hardness information accumulates over time, offering practical design guidance for real-time control loops.

Importantly, while the Cartesian robot provides the controlled setting required to isolate the role of load and speed, real manipulation involves fewer repeatable contacts and coupled hand dynamics. For this reason, the next chapter extends the hardness-discrimination objective to a prosthetic-hand platform (Hannes) in a less controlled environment, utilizing the methodological foundation developed here, which includes compact online features, incremental decision-making, and uncertainty characterization as a bridge between controlled bench-top validation and realistic grasping scenarios.

CHAPTER 4

Real-Time Hardness Classification Using Hannes hand and Piezoelectric Sensors

4.1 Introduction

In human manipulation, cutaneous mechanoreceptors provide high-bandwidth information about contact onset, pressure evolution, micro-vibrations, and local deformation; this information supports both reactive control (e.g., regulating grip force, preventing slip) and perceptual inference (e.g., judging compliance/hardness, texture, and contact quality). Replicating these capabilities in artificial hands has therefore been a central theme in tactile robotics and dexterous manipulation research, motivating the development of tactile skins and fingertip sensors, as well as learning-based perception pipelines that map tactile signals to object properties and manipulation-relevant states [187, **?**, 188].

Modern myoelectric prosthetic hands have reached a level of mechanical sophistication that enables stable grasps and human-like motion, yet their interaction with the environment remains fundamentally limited by the lack of rich, fast, and robust tactile perception at the fingertips. In human manipulation, cutaneous mechanoreceptors provide high-bandwidth information about contact onset, pressure evolution, micro-vibrations, and local deformation, while also conveying material-related cues that humans routinely exploit during manipulation (e.g., compliance and softness); this information supports both reactive control (e.g., regulating grip force, preventing slip) and perceptual inference (e.g., judging compliance/hardness, texture, and contact quality). In prosthetic use, this missing channel increases reliance on vision, reduces robustness under natural variability, and limits the user's ability to develop intuitive, low-effort strategies for object interaction [189, 190, 122]. From this perspective, enabling a prosthetic hand to extract interaction-relevant tactile descriptors is not merely a sensing add-on: it is a prerequisite for closing the loop toward practical assistance, whether through improved autonomous grasp regulation, shared control, or sensory feedback to the user.

Recent progress in tactile sensing and electronic skin has made this goal technically plausible, with demonstrations of high-density, flexible, and multimodal sensing architectures that can be integrated on anthropomorphic platforms [55, 191]. At the same time, the robotics and haptics communities have developed an extensive body of work on extracting object properties from tactile data, including approaches based on engineered descriptors as well as learning-based representations [28, 192]. However, a persistent challenge for prosthetic translation is that material-property inference is strongly affected by *how* contact occurs: variations in loading level, contact duration, and interaction speed alter both the amplitude and the temporal structure of tactile signals. As a result, models trained in tightly controlled conditions may degrade in performance when deployed under realistic grasp dynamics, where the same object can be contacted under multiple closing levels and multiple actuation speeds. This gap is particularly relevant for hardness/compliance discrimination, which is inherently confounded by the interaction trajectory and not only by the material itself [28, 189].

Motivated by this need to move hardness perception from controlled benchtop conditions toward prosthetic operation, this chapter extends the tactile-sensing activity of Chapter 3 to a less controlled yet application-relevant setting using the Hannes prosthetic hand. The focus is not on proposing a new neural network architecture per se, but on establishing a complete pipeline—sensorization, dataset design under realistic variability, offline model selection under latency constraints, and online deployment logic—that can support real-time hardness inference as a building block for future prosthetic functionalities. Concretely, we integrate a P(VDF-TrFE) piezoelectric tactile patch on the index-finger glove of Hannes and build a dataset that explicitly spans variability in grasp execution by combining five TPU objects of distinct infill percentages (hardness levels), three target closed positions (TCP) reflecting different squeezing levels, and four grasping frequencies (GF) reflecting different closing/opening speeds. We then study how temporal segmentation affects the accuracy–latency trade-off by training and comparing classical feature-based models against an end-to-end 1D-CNN across multiple non-overlapping window lengths. Beyond aggregate accuracy, we analyze scenario-wise performance across TCP–GF conditions and inspect the probabilistic outputs of the CNN through Top-k behavior and confidence calibration, which motivates confidence-aware

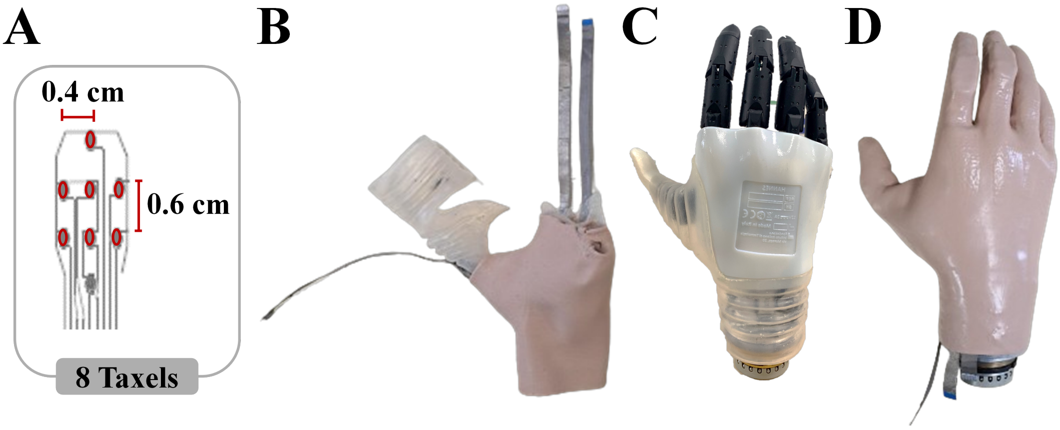

Figure 4.1: Piezoelectric sensing on the Hannes hand prosthesis: A) Tactile sensing patch, B) Fully integrated sensorized glove, C) Hannes hand with D) the sensorized glove.

decision rules for deployment [186]. Finally, we deploy the selected model in a real-time event-triggered pipeline on the prosthetic setup and quantify event-level accuracy together with the latency to correct high-confidence evidence, thereby validating whether offline performance transfers to online grasp events under natural variability.

This chapter is organized as follows. Section 4.2 describes the tactile sensing integration on the Hannes hand and the experimental setup used to generate a dataset under controlled variations of TCP and GF. The signal preprocessing, segmentation, feature extraction, model training protocol, and the real-time event-based pipeline are detailed in the Methods section number 4.3, with the deployed pipeline summarized in Section 4.3.5. Section 4.4 reports the offline model comparison across window lengths, the scenario-wise analysis, the probabilistic reliability results, and the online grasp evaluation. We then discuss how these findings relate to the state of the art in tactile property inference and prosthetic deployment in Section 4.5, and conclude in Section 4.6 by outlining limitations and directions toward using hardness estimates as a practical variable for prosthetic control and sensory feedback.

4.2 Materials

Following the developed activity in chapter 3, and to integrate the sensors in a real life scenario, a less controlled environment was developed to test the capabilities of the sensor in such environment. For that reason in this section we first present the sensorization of the Hannes hand using P(VDF-TrFE) piezoelectric sensors. Following, we present the object selection thought which we preliminary tested our setup. Finally, we describe the experimental setup and data collection.

4.2.1 Tactile Sensing integrated within hand glove

This study involved the two main functional blocks. The first one is represented by the piezoelectric sensing arrays integrated within the thumb, index, and middle fingertips' hand gloves, and the interface electronics (IE) to collect the voltage of the sensors. Sensorization of the index finger only was used for this experiment. The second block represents the host PC through which two GUI run in parallel: the first one is the HannesApp, used to control and collect the data of the hand, and the second one, a LabView GUI, used to log and analyze the data of the IE.

4.2.1.1 Sensing System: Overview

The proposed sensor technology was previously used in robotic applications [167][193]. Each sensing patch consists of 8 sensing units of 1 mm diameter each and 0.6 cm center-to-center pitch. What characterizes these sensors the most is their frequency bandwidth which ranges from 0.5 Hz to 1 KHz. The voltage output of the sensors is processed using an interface electronics (IE) equipped with an ARM Cortex M0 microcontroller and a DDC232 analog-to-digital converter, with a 2K samples/sec as a sampling rate. As a consequence, it was decided to integrate the sensing arrays on the Hannes hand to collect their response for targeting a real-use prosthetic application toward objects grasp. For this purpose, piezoelectric sensing patches were integrated, but only the one on the index finger was used in this experiment.

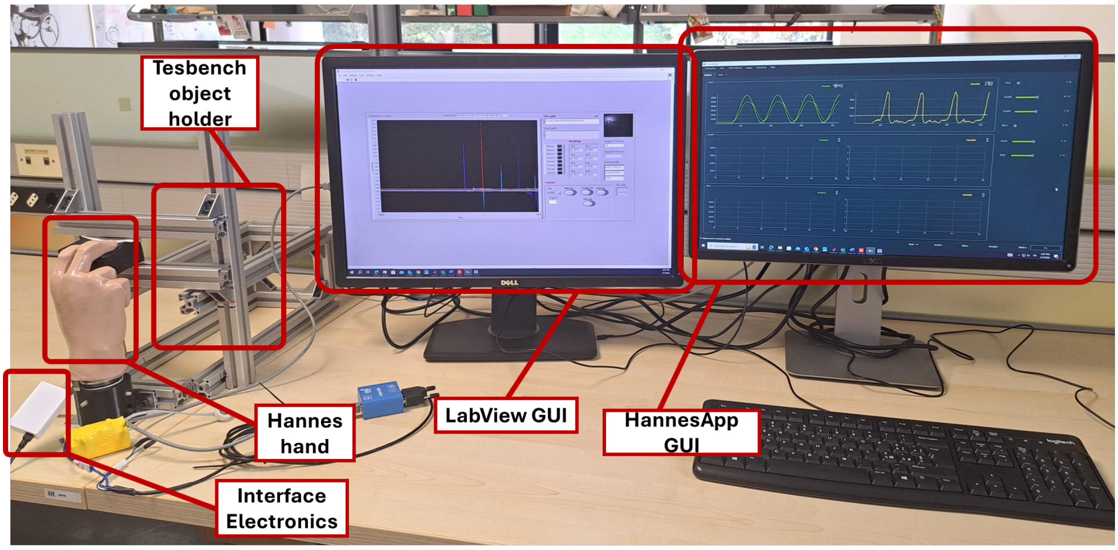

Figure 4.2: Experimental setup presenting the sensorized hand prosthesis and GUIs

4.2.1.2 Preparation of the sensing system on the prosthesis

We used the Hannes hand [33][194] because of its intrinsic human-like biomimetic performances and capabilities. Therefore, improving these hand functionalities could play a key role. A new integration strategy, distinct from the one previously reported in [40], was developed for the Hannes prosthetic hand. The sensing arrays were shielded and protected using the same materials and methods explained in section 3.2 where each sensing array was covered on both sides with conductive tape (Model tesa 60262, tesa) to reduce the sensitivity of the sensors against the noise, and external charges. The resulting structure was then sandwiched between the two elastic gloves of the Hannes hand using silicone rubber adhesive (Sil Proxy) (see Figure 4.1 B). As a result, removable sensors were obtained as illustrated in Figure 4.1 D). The IE presented in section 3.3 was also used in this system, and it was placed in a shielded box as shown in figure 4.2.

4.2.2 Experimental Setup and Data Collection

To collect experimental sensor data for hardness discrimination, 3D-printed objects presented and used in [166] were utilized in this study. Three cubic objects (7 x 7 x 4 cm), and two cylinder objects with a diameter of 4 cm and a height of 7 cm were 3D printed with five different filling percentages (3, 5, 7, 10, and 12%) using thermoplastic polyurethane (TPU). Figure 4.2 shows the experimental setup used to collect the experimental sensor data. Hannes prosthetic hand was used to perform repeated grasp actions on the 3D-printed objects while varying the grasping frequency and closed hand position. To do the data collection the prosthetic hand and the grasped object were fixed to ensure a consistent grasp, for that reason, a test bench was developed. The ad-hoc test bench was developed (Figure 4.2) composed of movable parallel arms and a two-cuff system acting as a holder to fix the objects in between. Hannes hand was fixed at the base of the test bench, in such way the thumb, index, and middle fingers had an impact on the objects when grasping. Hannes was connected to the PC through a CAN-USB adapter and controlled through HannesApp, which also allowed data acquisition (current consumption of the DC motor and the related encoder position). For the sensing side, the tactile sensors signals were collected through the IE which interfaced to the host PC through USB connection, then acquired and visualized through a graphical user interface (GUI) developed in NI LabVIEW on the host PC. In this study, grasping experiments were planned to perform human-like grasps to perceive the hardness of the objects in real-time while controlling the frequency of the grasp and the closing on the objects. To this end, a programmable mechanical input was applied by programming the hand to grasp the 3D objects using the sensorized fingertip of the index, while continuously measuring the electrical response of the PVDF sensors. The grasping procedure consisted of the following three steps at a programmed Grasping Frequency (GF): Open hand position, closing the hand until reaching the Target Closed Position (TCP), decreasing the load to the initial load, and then returning to the initial open position.

The grasping procedure was repeated for 3 TCP {1=75, 2=85, and 3=95} (%) at 4 GF {0.2, 0.3, 0.4, and 0.5} (Hz). Then repeated for each 3D-printed object, thus generating a dataset composed of 60 combinations of objects, TCP, and GF. Finally, each combination was repeated for 30 trials. The dataset can be formulated as

$$\mathcal{D} = \left\{ (X_i, y_i) \;\middle|\; \begin{array}{l} X_i \in \mathbb{R}^{N_C \times N_S}, \\ y_i \in \{\text{five classes}\}, \\ i = 1, \ldots, 1800 \end{array} \right\} \tag{4.1}$$

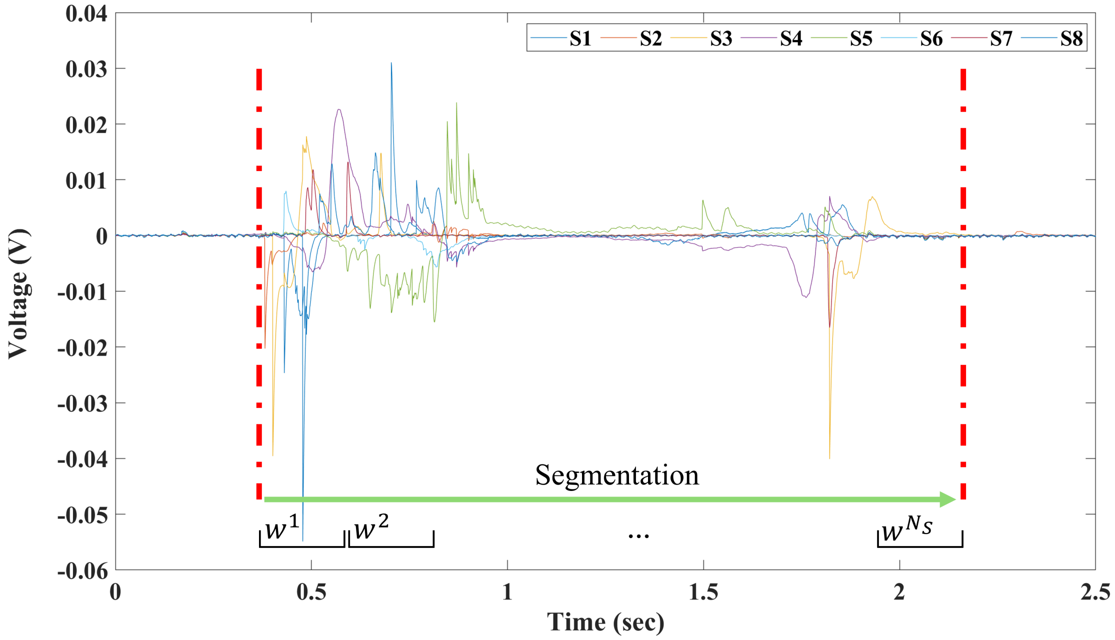

Figure 4.3: Figure shows segmentation of the signal during the response of the sensor without overlap between the windows

where X_i is a 2D tensor with $N_C = 8$ representing the number of channels (i.e., the sensors) and N_S the number of samples that depends on the TCP and GF as shown in Figure 4.3

4.3 Methods

In order to classify hardness using the developed tactile system in a less controlled environment than that of Chapter 3, and to build a real-time pipeline suitable for future prosthetic-hand applications, the following methodology was adopted. As illustrated in Fig. 4.4, the process starts with data preprocessing, where the signals are filtered and segmented into windows. These windows are then used as input to the feature and network selection block, in which classical machine-learning and deep-learning approaches are compared in order to identify the most suitable strategy for hardness classification with this sensing system. Finally, based on the offline training and selection stage, a real-time classification pipeline is implemented and optimized.

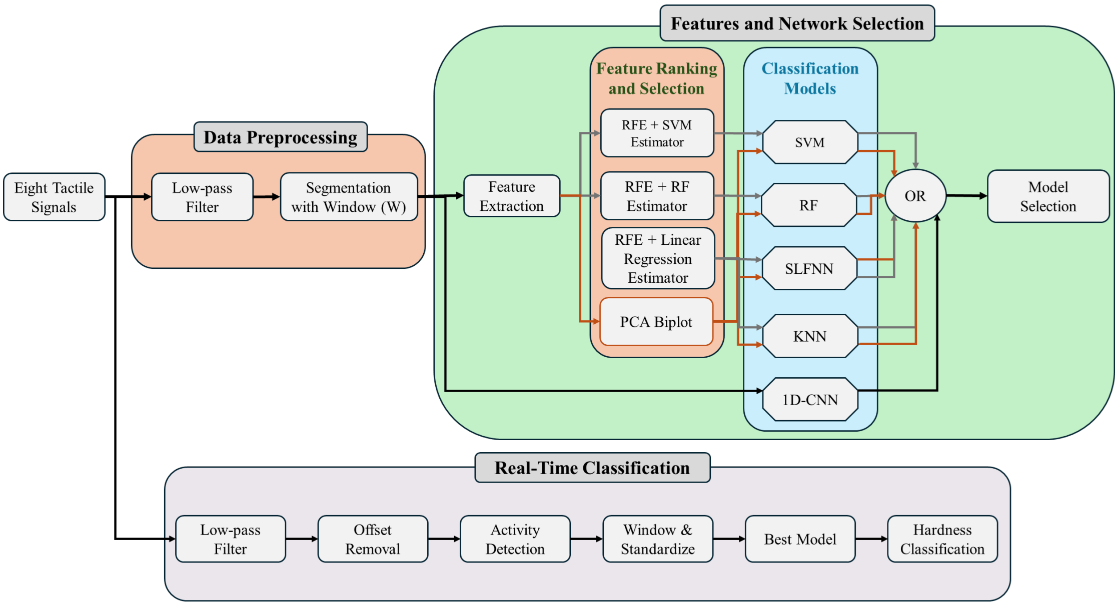

Figure 4.4: Block diagram explaining the proposed methodology to process the piezoelectric data, and build the real-time pipeline

4.3.1 Data Preprocessing

4.3.1.1 Signal Filtering

A 30 Hz low-pass filter was applied to each $X_i \in \mathcal{D}$ to attenuate high-frequency noise while retaining the informative components of the tactile signals, following the procedure in [171]. In addition, only the segments corresponding to the indentation phase were retained: the samples preceding the touch onset and those following the touch release (see Figure 4.3) were discarded. The resulting signals are denoted by $\widehat{X}_i \in \mathbb{R}^{N_C \times \widehat{N_S}}$, where $\widehat{N_S} < N_S$ reflects the removal of the non-touch-related portions.

4.3.1.2 Segmentation

In this study, the filtered signals $\widehat{X}_i \in \mathbb{R}^{N_C \times \widehat{N_S}}$ were segmented using fixed-size, non-overlapping windows defined in the time domain. Each trial was partitioned into contiguous windows of constant duration, and no overlap was introduced between consecutive segments.

Given the sampling frequency $f_s = 2$ kHz, a window of duration T_w milliseconds corresponds to

$$n_w(T_w) = T_w \cdot \frac{f_s}{1000} \tag{4.2}$$

samples. We explored multiple window durations T_w in the range $[20, 400]$ ms to study the trade-off between temporal resolution (short windows) and the amount of tactile information contained in each segment (long windows), and then how fast can we have a reliable information the system can use it to make a decision later on. For instance, at $f_s = 2$ kHz, $T_w = 20$ ms and $T_w = 400$ ms correspond to windows of $n_w = 40$ and $n_w = 800$ samples, respectively.

For a given window duration T_w, each filtered signal $\widehat{X}_i$ was partitioned into $M_i^{(T_w)}$ non-overlapping windows by grouping consecutive samples:

$$W_{i,j}^{(T_w)} = \widehat{X}_i\big[\,:,\ (j-1)\,n_w(T_w) + 1 : j\,n_w(T_w)\big], \quad j = 1, \ldots, M_i^{(T_w)}, \tag{4.3}$$

where $W_{i,j}^{(T_w)} \in \mathbb{R}^{N_C \times n_w(T_w)}$ denotes the j-th window extracted from trial i for the chosen duration T_w. The number of windows $M_i^{(T_w)}$ depends on the trial length and generally differs across trials.

The dataset associated with a given window duration T_w is then defined as

$$\mathcal{D}_{T_w} = \left\{ \left(W_{i,j}^{(T_w)}, y_i\right) \left| \begin{array}{l} W_{i,j}^{(T_w)} \in \mathbb{R}^{N_C \times n_w(T_w)}, \\ y_i \in \{\text{five classes}\}, \\ i = 1, \ldots, N, \\ j = 1, \ldots, M_i^{(T_w)} \end{array} \right. \right\}, \tag{4.4}$$

where N is the number of trials, N_C is the number of sensing channels, $n_w(T_w)$ is the number of samples per window determined by the duration T_w, $M_i^{(T_w)}$ is the number of windows obtained from trial i, and y_i is the hardness label associated with trial i. This formulation is repeated for each tested window duration $T_w \in [20, 400]$ ms to assess how the window length affects classification performance.

4.3.2 Feature extraction

To reduce the dimensionality of the input space while retaining the most informative descriptors of the tactile dynamics, a multi-stage feature extraction and selection pipeline is applied, as illustrated in Fig. 4.4. For each windowed trial $W_{i,j}^{(T_w)} \in \mathbb{R}^{N_C \times n_w(T_w)}$ (see Section 4.3.1.2), a set of time-domain features is computed independently on each channel and then concatenated into a single feature vector. This defines a mapping

$$z_{i,j} = \Phi\big(W_{i,j}^{(T_w)}\big) \in \mathbb{R}^D, \tag{4.5}$$

where $\Phi(\cdot)$ denotes the feature extraction operator and D is the total number of extracted features.

For each channel c, let $w_c[t]$ denote the t-th sample of the considered window, with

$$c = 1, \dots, N_C, \qquad t = 1, \dots, n_w(T_w).$$

Let

$$\mu_c = \frac{1}{n_w(T_w)} \sum_{t=1}^{n_w(T_w)} w_c[t]$$

denote the sample mean of channel c. The following 12 time-domain descriptors are then computed for each channel:

Mean

$$\mathrm{Mean}_c = \mu_c.$$

This feature represents the average signal level within the window.

Standard deviation

$$\mathrm{StD}_c = \sqrt{\frac{1}{n_w(T_w)} \sum_{t=1}^{n_w(T_w)} \big(w_c[t] - \mu_c\big)^2}.$$

It measures the dispersion of the tactile signal around its mean.

Minimum

$$\mathrm{Minimum}_c = \min_{1 \le t \le n_w(T_w)} w_c[t].$$

This captures the smallest amplitude observed in the window.

Maximum

$$\text{Maximum}_c = \max_{1 \le t \le n_w(T_w)} w_c[t].$$

This captures the largest amplitude observed in the window.

Median

$$\text{Median}_c = \text{median}\{w_c[1], \ldots, w_c[n_w(T_w)]\}.$$

The median provides a robust estimate of the central tendency of the signal.

Skewness

$$\text{Skewness}_c = \frac{1}{n_w(T_w)} \sum_{t=1}^{n_w(T_w)} \left(\frac{w_c[t] - \mu_c}{\text{StD}_c} \right)^3 .$$

Skewness quantifies the asymmetry of the signal distribution.

Kurtosis

$$\text{Kurtosis}_c = \frac{1}{n_w(T_w)} \sum_{t=1}^{n_w(T_w)} \left(\frac{w_c[t] - \mu_c}{\text{StD}_c} \right)^4 .$$

Kurtosis measures the peakedness or impulsiveness of the signal distribution.

Peak-to-peak amplitude

$$\text{PtP}_c = \text{Maximum}_c - \text{Minimum}_c.$$

This descriptor measures the dynamic range of the signal within the window.

Root mean square

$$\text{RMS}_c = \sqrt{\frac{1}{n_w(T_w)} \sum_{t=1}^{n_w(T_w)} w_c[t]^2}.$$

RMS reflects the effective energy of the signal.

Crest factor

$$\text{Crest}_c = \frac{\max_{1 \le t \le n_w(T_w)} |w_c[t]|}{\text{RMS}_c}.$$

Crest factor quantifies the prominence of peaks relative to the average signal energy.

Zero-crossing rate

$$\mathrm{ZCR}_c = \frac{1}{n_w(T_w) - 1} \sum_{t=1}^{n_w(T_w)-1} \mathbb{I}(w_c[t]\, w_c[t+1] < 0)\,,$$

where $\mathbb{I}(\cdot)$ is the indicator function. This feature measures how frequently the signal changes sign and is therefore sensitive to local oscillatory behaviour.

Waveform length

$$\mathrm{WFL}_c = \sum_{t=1}^{n_w(T_w)-1} |w_c[t+1] - w_c[t]|\,.$$

Waveform length captures the cumulative amount of variation in the signal and is related to its temporal complexity.

The 12-dimensional feature vector extracted from channel c is therefore

$$\boldsymbol{\phi}_c\left(W_{i,j}^{(T_w)}\right) = \begin{bmatrix} \mathrm{Mean}_c & \mathrm{StD}_c & \mathrm{Minimum}_c & \mathrm{Maximum}_c & \mathrm{Median}_c & \mathrm{Skewness}_c \\ \mathrm{Kurtosis}_c & \mathrm{PtP}_c & \mathrm{RMS}_c & \mathrm{Crest}_c & \mathrm{ZCR}_c & \mathrm{WFL}_c \end{bmatrix}^\top \in \mathbb{R}^{12}.$$

Given N_C sensing channels, the complete feature vector is obtained by concatenating the per-channel descriptors:

$$z_{i,j} = \begin{bmatrix} \boldsymbol{\phi}_1^\top & \boldsymbol{\phi}_2^\top & \cdots & \boldsymbol{\phi}_{N_C}^\top \end{bmatrix}^\top \in \mathbb{R}^{12N_C}.$$

Hence, $D = 12N_C$ features are extracted from each window. These descriptors include basic statistics, higher-order moments, amplitude-related measures, and shape descriptors, which jointly capture both the intensity and the temporal variability of the tactile signals. Most of them can be computed from running sums, extrema, or local sample differences, making them suitable for efficient implementation in embedded systems.

Before applying feature ranking methods, two unsupervised filtering steps are performed to remove uninformative or redundant descriptors. First, a variance thresholding step discards features whose variance across the training set is below a small threshold, thereby eliminating nearly constant features. Second, a correlation-based pruning step removes highly collinear descriptors by discarding one feature from each pair whose absolute Pearson correlation exceeds a predefined threshold (e.g., $|\rho| > 0.95$). The remaining features constitute a reduced set on which the subsequent ranking and selection procedures are applied.

For each window duration T_w, all feature filtering and ranking operations were fitted exclusively on the training split. In particular, the variance threshold, the correlation-pruning mask, and the subsequent PCA-based or RFE-based rankings were estimated using only the training feature matrix. The resulting feature masks and ordered subsets were then applied unchanged to the validation and test sets. This procedure was adopted to prevent information leakage from the validation or test data into the feature-selection stage and to ensure that the reported performance reflects generalization to unseen samples [178, 179, 195].

4.3.2.1 PCA-based ranking

As in previous work [173], principal component analysis (PCA) [172] is employed as an unsupervised tool for feature ranking. PCA transforms the standardized feature matrix into a new set of orthogonal directions, the principal components (PCs), which are linear combinations of the original variables and ordered according to the amount of variance they explain.

In this study, PCA is applied after variance and correlation pruning. The remaining features are first standardized, and PCA is then computed with two components (PC1 and PC2). Let $\ell_{j,1}$ and $\ell_{j,2}$ denote the loadings of feature j on PC1 and PC2, respectively. An importance score is assigned to each feature based on the magnitude of its contributions to these two components:

$$s_j = |\ell_{j,1}| + |\ell_{j,2}|. \tag{4.6}$$

Features are then sorted in descending order of s_j, yielding a global ranking in which higher-scoring descriptors are considered more relevant. From this ordered list, the top K features (with $K = 10$ in this work) are selected and used to train and tune the classical machine learning classifiers (SVM, KNN, RF, and SLFNN). This PCA-based ranking provides an unsupervised, model-agnostic criterion that highlights features most aligned with the directions of maximum variance in the data.

4.3.2.2 RFE-based ranking

Recursive feature elimination (RFE) [174] is adopted as a supervised feature selection method, applied on the pruned feature set to obtain model-specific rankings. RFE iteratively fits a predictive model, evaluates feature importance, and removes the least relevant descriptors until a target number of features is reached. This process yields both a ranked list of features and a reduced subset tailored to the chosen estimator.

In the proposed pipeline, RFE is performed separately for each classifier type. For the SVM and RF configurations, linear support vector machines and random forests, respectively, are used as base estimators within the RFE procedure, leveraging their native importance measures (model coefficients for linear SVMs and impurity-based importances for RFs). For classifiers that do not directly provide suitable feature importance scores (e.g., KNN and SLFNN), a linear model with sparsity-promoting regularization is used within RFE to estimate feature relevance. In all cases, the number of selected features is fixed to $K = 10$, matching the dimensionality used in the PCA-based selection.

Formally, let $\mathcal{F} = \{f_1, \ldots, f_M\}$ denote the set of features remaining after variance and correlation pruning. RFE produces, for each model m, an ordered sequence

$$\mathcal{F}_{\text{RFE}}^{(m)} = \big(f_{\pi^{(m)}(1)}, \ldots, f_{\pi^{(m)}(M)}\big), \tag{4.7}$$

where $\pi^{(m)}(\cdot)$ is a permutation that ranks features from most to least relevant according to estimator m. The first K elements of $\mathcal{F}_{\text{RFE}}^{(m)}$ define the model-specific subset used to train and tune the corresponding classifier. This supervised selection strategy complements the PCA-based ranking by explicitly accounting for the discriminative power of each feature with respect to the hardness classification task.

4.3.3 Model selection and training

In this study, we compare a deep learning model operating directly on the raw windowed signals $W_{i,j}^{(T_w)} \in \mathbb{R}^{N_C \times n_w(T_w)}$ with a set of classical machine learning models trained on the corresponding feature vectors $z_{i,j} = \Phi\big(W_{i,j}^{(T_w)}\big)$ described in the previous section.

The goal is to evaluate the trade-off between representation learning from raw data and classification based on compact, hand-crafted descriptors, under constraints that remain compatible with potential embedded deployment.

A one-dimensional convolutional neural network (1D-CNN) is adopted as the deep learning baseline, as it can effectively capture local temporal patterns from fixed-length windows of the tactile signal. In parallel, four classical models are considered: Support Vector Machines (SVM), k-Nearest Neighbours (KNN), Random Forests (RF), and a Single-Layer Feedforward Neural Network (SLFNN). These models provide lighter-weight alternatives that operate in the feature domain and offer varying degrees of interpretability and flexibility.

All models are trained independently for each window duration T_w (see Section 4.3.1.2). For classical machine learning, the input consists of the feature vectors $z_{i,j}$ obtained after extraction, variance filtering, correlation pruning, and feature selection (either PCA-based or RFE-based ranking). For the CNN, the input is the normalized multichannel time series window $W_{i,j}^{(T_w)}$. In all cases, the data are partitioned into three mutually exclusive sets: training, validation, and test. Hyperparameters are tuned using the validation set, and final performance is assessed on the held-out test set using the configuration that maximizes validation accuracy.

4.3.3.1 Data splitting, hyperparameter tuning, and model-selection protocol

To ensure a fair comparison across models and window durations, the data partitioning was performed at the trial level before segmentation. Thus, all windows extracted from a given trial belong to the same subset, and no windows originating from the same physical interaction appear simultaneously in training, validation, and test. For each tested duration T_w, the trial-wise training, validation, and test partitions were segmented independently, yielding three mutually exclusive windowed datasets associated with that temporal resolution.

For the classical machine-learning branch, the complete pipeline was executed independently for each pair (T_w, feature-selection strategy). First, feature extraction was applied to the windows of the three subsets. Then, variance filtering, correlation pruning, and feature ranking/selection were fitted on the training subset only and transferred unchanged to the validation and test subsets. For every candidate model and hyperparameter configuration, the classifier was trained on the training subset and evaluated on the validation subset. The validation accuracy was used as the model-selection criterion. After selecting the best hyperparameter configuration for a given classical pipeline, the corresponding classifier was refitted using the available development data and then evaluated once on the held-out test subset.

For the deep-learning branch, the 1D-CNN was trained independently for each window duration T_w and for each candidate hyperparameter configuration. Model selection was based on the validation subset, whereas the test subset was kept untouched until the final evaluation of the selected configuration. This train/validation/test strategy was adopted in order to separate hyperparameter optimization from final performance assessment and to avoid optimistic bias in the reported results [178, 179, 195].

4.3.3.2 Convolutional Neural Network (1D-CNN)

The 1D-CNN processes windows $W_{i,j}^{(T_w)} \in \mathbb{R}^{N_C \times n_w(T_w)}$ directly in the time domain. Prior to training, each channel is standardized using the mean and standard deviation computed over the training set, and the same normalization is applied to the validation and test sets.

The network is composed of a sequence of convolutional blocks followed by a fully connected classifier. Each block consists of a 1D convolutional layer with ReLU activation, optionally followed by a max-pooling operation along the temporal dimension to progressively reduce signal length while increasing the number of feature maps. The number of convolutional blocks and the base number of filters are treated as hyperparameters: specifically, the number of filters in the first convolutional layer is set to either 16 or 32, deeper blocks use progressively larger channel counts, the kernel size is chosen between 3 and 5 samples, and the number of convolutional blocks is set to either 3 or 7. The fully connected part consists of a dense layer with 128 neurons and ReLU activation followed by a dropout layer, where the dropout rate is selected between 0.2 and 0.3, and a final softmax output layer with one neuron per class. The network is trained using the Adam optimizer with a learning rate fixed to 10^{-3} and a cross-entropy loss.

For each hyperparameter configuration, the model is trained with an early stopping criterion based on the validation loss, using a maximum of 100 epochs and a patience of 5 epochs. The weights corresponding to the lowest validation loss are retained. Among all tested configurations, the one yielding the highest validation accuracy is selected, and the corresponding model is finally evaluated on the test set.

4.3.3.3 Support Vector Machine (SVM)

SVMs are trained on the selected feature subsets obtained from the PCA-based and RFE-based rankings applied to the vectors $z_{i,j}$. A radial basis function (RBF) kernel is adopted to capture nonlinear relationships in the feature space. The penalty parameter C is explored in the set $\{0.1, 1, 10\}$, while the kernel coefficient γ is tested using the `scale` and `auto` options provided by *scikit-learn*, which adapt the kernel width to the feature distribution. Multi-class classification is handled using the one-vs-one strategy. Features are standardized within a pipeline before model fitting, and the hyperparameter combination that maximizes validation accuracy is selected.

4.3.3.4 K-Nearest Neighbors (KNN)

The KNN classifier provides a non-parametric baseline that classifies each sample based on the labels of its nearest neighbors in the feature space spanned by $z_{i,j}$. The number of neighbors k is varied in the set $\{3, 5, 7, 9\}$, and two weighting schemes are compared: uniform weighting, where all neighbors contribute equally, and distance-based weighting, where closer neighbors have a higher influence. The Minkowski distance order p is set either to 1 (Manhattan distance) or 2 (Euclidean distance). As in the other classical models, features are standardized via a preprocessing step in the pipeline, and the configuration achieving the highest validation accuracy is retained.

4.3.3.5 Random Forest (RF)

Random Forests constitute an ensemble of decision trees trained on bootstrapped samples of the data with feature-level randomness at each split. This model is particularly suitable for handling heterogeneous and potentially correlated features, and it naturally provides measures of feature importance.

To control the complexity and expressiveness of the ensemble, several hyperparameters are tuned. The number of trees $n_{\text{estimators}}$ is set to either 100 or 200, the maximum tree depth is chosen among $\{5, 6, 7\}$, the minimum number of samples required to split an internal node is set to either 2 or 5, and the minimum number of samples required to form a leaf node is set to either 1 or 2. For each window duration and feature selection method, the RF is fitted on the training set, evaluated on the validation set for hyperparameter selection, and finally assessed on the test set using the best-performing configuration.

4.3.3.6 Single Layer Feedforward Neural Network (SLFNN)

The SLFNN is implemented as a multilayer perceptron with a single hidden layer, trained on the selected feature sets $z_{i,j}$. This architecture provides a flexible yet relatively lightweight nonlinear classifier.

The size of the hidden layer is explored in the set $\{50, 100, 200\}$ neurons, with ReLU used as the activation function. To regularize the model and mitigate overfitting, the L2 penalty coefficient α is varied between 10^{-4} and 10^{-3}. Two learning rate schedules are considered, namely a constant learning rate and an adaptive scheme, both implemented within the *scikit-learn* MLP framework. Training is performed with early stopping based on an internal validation split and a maximum of 1000 iterations. As with the other models, the external validation set is used to select the optimal hyperparameter combination within a standardized pipeline.

4.3.3.7 Training protocol

For each window duration T_w and for each feature selection strategy (PCA-based ranking and RFE-based ranking), the SVM, KNN, RF, and SLFNN models are independently trained and tuned using the training and validation sets, operating on the feature vectors $z_{i,j}$. Feature standardization is performed within the model pipelines to avoid data leakage. The hyperparameter configuration that maximizes validation accuracy is retained, and the corresponding model is finally evaluated on the held-out test set. The 1D-CNN is trained separately on the raw normalized signals $W_{i,j}^{(T_w)}$ for each window duration, following the same train–validation–test protocol. This design enables a consistent comparison between deep and classical models across different temporal resolutions of the tactile signal.

4.3.4 Statistical analysis

To complement the descriptive evaluation based on classification accuracy, formal statistical analyses were performed in order to assess whether the observed performance differences were statistically meaningful. Since the compared classifiers were evaluated on the same held-out test samples, paired statistical procedures were adopted.

First, to evaluate *model selection* among the classical machine-learning approaches, the predictions obtained at the selected window size of 100 ms were compared across the complete classical pipelines retained in this chapter. Since feature selection was implemented in a classifier-specific manner, namely PCA-ranked subsets and RFE subsets built with different estimators depending on the classifier, inferential comparisons were performed at the level of full pipelines rather than through a single pooled feature-selection factor. Accordingly, the comparison involved SVM trained on $\mathcal{F}_{\text{100-RFE-SVM}}$, RF trained on $\mathcal{F}_{\text{100-RFE-RF}}$, and KNN and SLFNN trained on $\mathcal{F}_{\text{100-RFE-L1}}$. Because more than two matched pipelines were compared on the same test samples, Cochran's Q test was used as an omnibus test. When the omnibus test indicated a significant difference, post-hoc pairwise comparisons were carried out using McNemar's test with Holm correction for multiple comparisons [181, 180, 182].

Second, to support the main offline comparison between the best classical pipeline and the deep-learning approach, the RF classifier trained on $\mathcal{F}_{\text{100-RFE-RF}}$ was compared against the 1D-CNN at 100 ms using McNemar's test. This test was selected because it is specifically suited to paired comparisons between two classifiers evaluated on the same test set [181].

Finally, the influence of the experimental conditions on the final deep-learning model was analyzed at the trial level. Since the CNN generated one prediction per signal window, whereas the experimental conditions (TCP and GF) were defined for the original trial, window-level outputs belonging to the same original trial were first aggregated into a single trial-level decision. Specifically, the confidence vectors of all windows associated with the same trial were averaged, and the final predicted label for that trial was assigned as the class with the highest mean confidence. This soft-voting strategy was adopted in order to preserve the confidence information of the classifier while converting segmented predictions into a single decision per original trial, which is consistent with probability-aware handling of segmented time-series outputs [183]. After aggregation, trial-level correctness was analyzed as a function of TCP level, GF level, and their interaction (TCP × GF). Statistical significance was assessed at $p < 0.05$, and multiple pairwise comparisons were corrected using the Holm procedure whenever applicable.

4.3.5 Real-time classification pipeline

To enable online hardness discrimination, the predictive model selected in the offline analysis was deployed within a real-time pipeline that processes streaming tactile data from the sensing system. The real-time loop mirrors the preprocessing and windowing used during training, while adding activity detection, event-based triggering, and temporal smoothing of predictions.

The real-time classifier corresponds to the configuration selected during the offline train/validation/test procedure, including the window duration, the learned model parameters, and the normalization statistics estimated on the training set. Once selected offline, the model weights were kept fixed during real-time operation, and no online adaptation or incremental retraining was performed. This design was chosen to ensure that the real-time evaluation reflects the direct deployment of the previously validated offline model under streaming conditions.

Tactile data are acquired from the piezoelectric sensing array through a serial interface. Each frame consists of N_C channels encoded as 16-bit integers and is converted to physical units using a fixed scaling factor. The samples are streamed at a known sampling frequency f_s, and a deployment configuration specifies the number of channels, the sampling rate, the window duration T_w (or equivalently the number of samples $N_w = T_w f_s$), as well as the parameters of the real-time filters and the activity detector.

For each incoming sample, a smoothing stage is applied independently to every channel. This is implemented as an exponential moving average of the form

$$y_c(t) = \alpha x_c(t) + (1 - \alpha) y_c(t-1),$$

where α is chosen to realize a desired low-pass cutoff frequency. The filter acts as a low-pass stage that attenuates high-frequency noise while preserving the relevant tactile dynamics.

To compensate for slow drifts and quasi-static offsets, a baseline adaptation mechanism is applied after smoothing. A baseline $\mu_c(t)$ is estimated per channel using an exponential update with time constant τ, and the baseline-corrected signal is obtained as

$$x_c^{\mathrm{bc}}(t) = y_c(t) - \mu_c(t).$$

The baseline is updated only while the system is inactive (no contact), so that indentation-related variations do not contaminate the reference level.

On top of the baseline-corrected signal, an activity detector identifies contact events in real time. At each time step, an aggregate magnitude $m(t)$ is computed across channels, as the mean absolute value (L1 aggregation). During inactive periods, the detector maintains an online estimate of the mean μ_0 and variance σ_0^2 of $m(t)$ using incremental updates. An adaptive threshold

$$\theta(t) = \mu_0(t) + k_\sigma\, \sigma_0(t)$$

is then used to decide whether the current magnitude indicates an active contact. A finite-state machine with hysteresis enforces minimum activation and deactivation times: the system enters the active state only if $m(t) > \theta(t)$ for at least a predefined duration $T_{\mathrm{min,active}}$, and returns to the inactive state only if $m(t) \leq \theta(t)$ for at least $T_{\mathrm{min,inactive}}$. Each continuous active interval defines an *event*, corresponding to a single grasp or indentation episode. At event onset, the internal buffers are reset, the event start time is stored, and baseline adaptation is temporarily frozen until the event ends.

While the detector is in the active state, baseline-corrected samples $x_c^{\mathrm{bc}}(t)$ are accumulated in a buffer of length N_w, corresponding to the window duration T_w used during training. Once N_w consecutive samples have been collected, a window $W^{\mathrm{RT}} \in \mathbb{R}^{N_C \times N_w}$ is formed. This window is then normalized using the same channel-wise means and standard deviations estimated on the training set and stored together with the model parameters. The normalized window is arranged to match the input layout expected by the deployed model and passed through the inference engine running on the target device.

For each window, the model produces a vector of class scores (logits) $\boldsymbol{\ell} \in \mathbb{R}^{N_{\text{classes}}}$, which are converted into class probabilities via a softmax function:

$$p_k = \frac{\exp(\ell_k)}{\sum_{j=1}^{N_{\text{classes}}} \exp(\ell_j)}, \quad k = 1, \dots, N_{\text{classes}}.$$

The predicted class for that window is given by the index of the maximum probability, and the corresponding value $p_{\max}$ is retained as a confidence measure. To reduce prediction jitter across consecutive windows, a temporal voting scheme is applied: the last K window-level predictions within the event are stored, and the real-time label at each step is obtained as the majority class among these K predictions. This majority-vote label is reported together with the per-window softmax probabilities.

The real-time pipeline operates in an event-triggered mode, in which predictions are produced only while the activity detector signals an ongoing contact. For each event, all window-level predictions and their probabilities are recorded, and at the end of the event a summary is computed by taking the majority label across its windows and the fraction of windows supporting this label. During the experiments, a graphical interface was used to visualize the predicted label as a function of time together with the top-K class probabilities, enabling real-time monitoring of the system behaviour under different grasping conditions.

This implementation ensures that the deployed model uses the same preprocessing (smoothing, baseline correction, windowing, and normalization) and window length as in the offline training phase, while adding adaptive activity detection and temporal smoothing tailored to online operation.

4.4 Results

4.4.1 Feature Selection

For each fixed window duration between 50 ms and 200 ms, the PCA-based ranking and the three RFE configurations consistently yielded compact 10-dimensional subsets, To compactly denote each 10-dimensional subset, we use the notation $\mathcal{F}_{T_w\text{-}Y}$, where $T_w \in \{50, 60, 70, 80, 90, 100, 200\}$ ms indicates the window duration and $Y \in \{\text{PCA}, \text{RFE-SVM}, \text{RFE-RF}, \text{RFE-L1}\}$ denotes the feature selection method. Individual features are annotated as feature_j, where feature $\in$ $\{\texttt{mean}, \texttt{std}, \texttt{min}, \texttt{max}, \texttt{median}, \texttt{skew}, \texttt{kurtosis}, \texttt{ptp}, \texttt{rms}, \texttt{crest}, \texttt{zcr}, \texttt{wfl}\}$ denotes the descriptor type and $j \in \{0, \ldots, 7\}$ is the index of the sensor channel from which that descriptor was computed. For example, `min_0` denotes the minimum value of channel 0, `ptp_3` the peak-to-peak amplitude of channel 3, and `wfl_6` the waveform length of channel 6. The resulting 10-dimensional subsets are summarized in Tables 4.1–4.4.

The feature labels reported in Tables 4.1–4.4 therefore correspond directly to the formally defined descriptors introduced in the feature-extraction methodology.

Overall, the PCA-based subsets are dominated by amplitude- and envelope-related descriptors on a small group of sensors, such as *min_0*, *max_0*, *max_3*, and the pair *min_5*/*max_5*. Crest factor and, to a lesser extent, skewness also appear among the leading components, especially on sensors 0, 3, and 6, indicating that PCA tends to emphasise large-scale intensity variations and peak prominence in the tactile response.

In contrast, the RFE-based subsets are strongly estimator-dependent. When RFE is driven by an SVM (Table 4.2), the selected features are almost exclusively waveform length (*wfl_j*) and peak-to-peak values (*ptp_j*) across several sensors, with zero-crossing rate (*zcr_j*) appearing for longer windows (e.g., *zcr_6* and *zcr_7* at 200 ms). This pattern suggests that the SVM relies heavily on the overall temporal complexity and amplitude excursions of the signal rather than on absolute levels.

RFE with a Random Forest estimator (Table 4.3) preserves a similar emphasis on dynamic descriptors, repeatedly selecting *ptp_j* and *wfl_j* across sensors 0–7, while also retaining *zcr_1* and *zcr_7* as informative indicators of high-frequency activity. For the longest window, *min_4* is added, indicating that the Random Forest model benefits from combining local amplitude statistics with measures of signal roughness.

Finally, the subsets produced by RFE with an L1-regularised logistic regression estimator (Table 4.4)—used to generate the ranked features for both KNN and SLFNN—tend to prioritise shape-related descriptors such as crest factor, skewness, and zero-crossing rate. These features appear consistently on sensors 0, 1, 3, 4, 6, and 7, and for $T_w = 200$ ms, *kurtosis_7* also enters the subset. This behaviour highlights the sensitivity of the L1-based RFE to higher-order distributional properties and rare large peaks. Taken together, these results show that hardness information is encoded both in

Table 4.1: Selected 10-dimensional feature subsets obtained from PCA-based ranking for window durations between 50 ms and 200 ms.

Subset	Selected features
$\mathcal{F}_{50\text{-PCA}}$	min_0, max_3, min_3, max_0, min_5, max_5, crest_3, ptp_3, wfl_3, crest_0
$\mathcal{F}_{60\text{-PCA}}$	min_0, min_3, max_3, max_0, min_5, max_5, crest_3, crest_0, wfl_4, max_6
$\mathcal{F}_{70\text{-PCA}}$	max_0, min_0, min_3, max_3, crest_0, crest_3, skew_0, min_5, max_5, ptp_0
$\mathcal{F}_{80\text{-PCA}}$	min_0, max_3, max_0, min_3, min_5, max_5, ptp_3, crest_0, wfl_4, ptp_0
$\mathcal{F}_{90\text{-PCA}}$	min_0, max_0, max_3, min_3, min_5, max_5, crest_3, crest_0, ptp_3, wfl_3
$\mathcal{F}_{100\text{-PCA}}$	max_3, min_0, max_0, min_3, max_5, min_5, crest_3, ptp_3, wfl_3, wfl_4
$\mathcal{F}_{200\text{-PCA}}$	max_0, min_0, min_3, max_3, crest_0, skew_0, crest_3, min_5, max_5, max_6

Table 4.2: Selected 10-dimensional feature subsets obtained from RFE with an SVM estimator for window durations between 50 ms and 200 ms.

Subset	Selected features
$\mathcal{F}_{50\text{-RFE-SVM}}$	wfl_0, wfl_1, ptp_2, wfl_2, wfl_3, wfl_4, ptp_5, wfl_6, ptp_7, wfl_7
$\mathcal{F}_{60\text{-RFE-SVM}}$	wfl_1, ptp_2, wfl_2, wfl_3, wfl_4, ptp_5, zcr_6, wfl_6, ptp_7, wfl_7
$\mathcal{F}_{70\text{-RFE-SVM}}$	wfl_0, wfl_1, ptp_2, wfl_2, wfl_3, wfl_4, ptp_5, wfl_6, ptp_7, wfl_7
$\mathcal{F}_{80\text{-RFE-SVM}}$	wfl_1, ptp_2, wfl_2, wfl_3, wfl_4, ptp_5, zcr_6, wfl_6, ptp_7, wfl_7
$\mathcal{F}_{90\text{-RFE-SVM}}$	wfl_1, ptp_2, wfl_2, wfl_3, wfl_4, ptp_5, zcr_6, wfl_6, ptp_7, wfl_7
$\mathcal{F}_{100\text{-RFE-SVM}}$	wfl_1, ptp_2, wfl_2, wfl_3, wfl_4, ptp_5, zcr_6, wfl_6, ptp_7, wfl_7
$\mathcal{F}_{200\text{-RFE-SVM}}$	wfl_0, wfl_1, wfl_2, zcr_4, wfl_4, zcr_5, zcr_6, wfl_6, zcr_7, wfl_7

the global amplitude of the tactile response (captured by PCA-selected features) and in its temporal structure and asymmetry (captured by the different RFE configurations).

Table 4.3: Selected 10-dimensional feature subsets obtained from RFE with a Random Forest estimator for window durations between 50 ms and 200 ms.

Subset	Selected features
$\mathcal{F}_{50\text{-RFE-RF}}$	min_0, ptp_0, wfl_1, ptp_3, wfl_3, wfl_4, ptp_5, wfl_5, ptp_6, wfl_6
$\mathcal{F}_{60\text{-RFE-RF}}$	ptp_0, ptp_3, wfl_3, ptp_4, wfl_4, ptp_5, wfl_5, ptp_6, wfl_6, zcr_7
$\mathcal{F}_{70\text{-RFE-RF}}$	ptp_0, zcr_1, ptp_3, wfl_3, ptp_4, wfl_4, ptp_5, wfl_5, wfl_6, zcr_7
$\mathcal{F}_{80\text{-RFE-RF}}$	ptp_0, zcr_1, ptp_3, wfl_3, ptp_4, wfl_4, ptp_5, wfl_5, wfl_6, zcr_7
$\mathcal{F}_{90\text{-RFE-RF}}$	ptp_0, zcr_1, ptp_3, ptp_4, wfl_4, ptp_5, wfl_5, wfl_6, ptp_7, zcr_7
$\mathcal{F}_{100\text{-RFE-RF}}$	ptp_0, zcr_1, ptp_3, wfl_3, ptp_4, wfl_4, ptp_5, wfl_5, wfl_6, zcr_7
$\mathcal{F}_{200\text{-RFE-RF}}$	ptp_0, zcr_1, ptp_3, min_4, ptp_4, wfl_4, wfl_5, wfl_6, ptp_7, zcr_7

Table 4.4: Selected 10-dimensional feature subsets obtained from RFE with an L1-regularised Logistic Regression estimator (used to rank features for both KNN and SLFNN) for window durations between 50 ms and 200 ms.

Subset	Selected features
$\mathcal{F}_{50\text{-RFE-L1}}$	zcr_0, zcr_1, zcr_3, crest_4, zcr_5, crest_6, zcr_6, skew_7, crest_7, zcr_7
$\mathcal{F}_{60\text{-RFE-L1}}$	zcr_0, zcr_1, crest_4, zcr_5, skew_6, crest_6, zcr_6, skew_7, crest_7, zcr_7
$\mathcal{F}_{70\text{-RFE-L1}}$	zcr_0, crest_1, crest_4, skew_5, crest_5, skew_6, crest_6, skew_7, crest_7, zcr_7
$\mathcal{F}_{80\text{-RFE-L1}}$	crest_0, crest_1, crest_3, crest_4, skew_5, crest_5, skew_6, crest_6, skew_7, crest_7
$\mathcal{F}_{90\text{-RFE-L1}}$	zcr_0, zcr_1, crest_4, zcr_5, skew_6, crest_6, zcr_6, skew_7, crest_7, zcr_7
$\mathcal{F}_{100\text{-RFE-L1}}$	crest_0, crest_1, crest_3, crest_4, skew_5, crest_5, skew_6, crest_6, skew_7, crest_7
$\mathcal{F}_{200\text{-RFE-L1}}$	skew_0, crest_0, crest_1, crest_3, crest_4, skew_5, crest_5, crest_6, kurtosis_7, crest_7

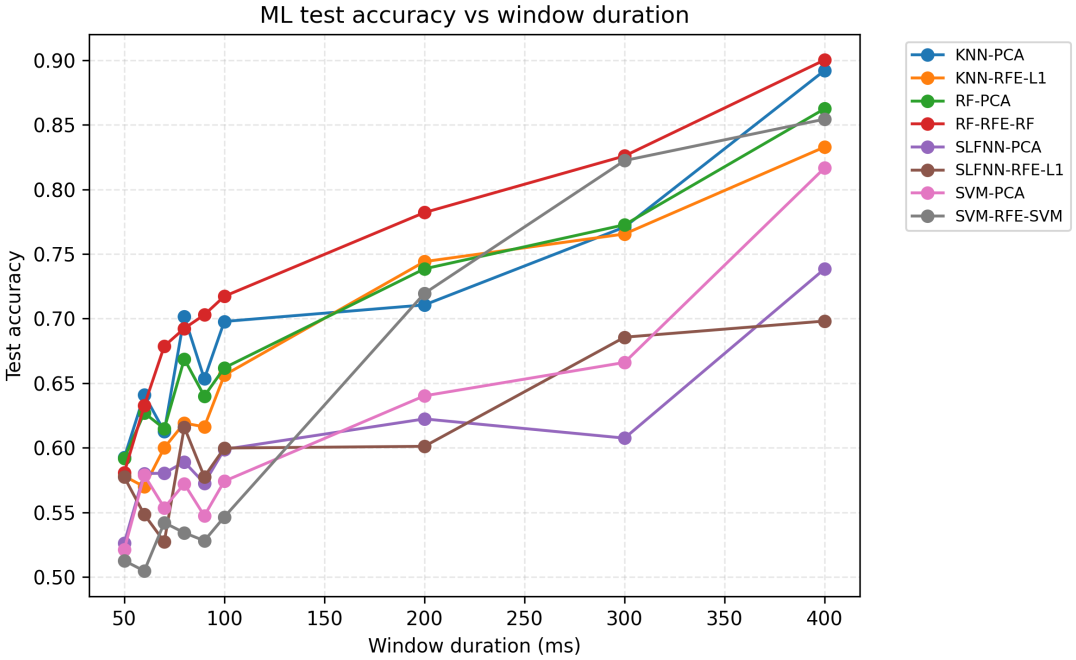

Figure 4.5: Test accuracy of all machine learning models as a function of window duration T_w. Each curve corresponds to a specific classifier and feature–selection method (KNN, RF, SLFNN, SVM combined with PCA or RFE-based subsets).

4.4.2 Classification Performance and Model Selection

For each window duration $T_w \in \{50, 60, 70, 80, 90, 100, 200, 300, 400\}$ ms, all models were trained on the training set, tuned on the validation set, and finally evaluated on a held-out test set. The four feature-based classifiers (KNN, RF, SLFNN, and SVM) were combined with the feature subsets introduced in Section 4.4.1 (PCA-ranked features $\mathcal{F}_{T_w\text{-PCA}}$, RFE-SVM subsets $\mathcal{F}_{T_w\text{-RFE-SVM}}$, RFE-RF subsets $\mathcal{F}_{T_w\text{-RFE-RF}}$, and RFE-L1 subsets $\mathcal{F}_{T_w\text{-RFE-L1}}$). The resulting test accuracies are summarised in Fig. 4.5.

As shown in Fig. 4.5, increasing the window duration systematically improves the performance of all feature-based models. For very short windows ($T_w = 50$ ms), test accuracies span approximately 50.5%–59.2% across all classifier–feature combinations, with the best configuration being KNN trained on the PCA subset $\mathcal{F}_{50\text{-PCA}}$ (test accuracy 59.2%). When the window is extended to $T_w = 60$ ms, the best ML accuracy rises to 64.1%, again obtained with KNN and PCA-ranked features. For intermediate windows ($T_w = 70$–100 ms), the top-performing configurations alternate between KNN with

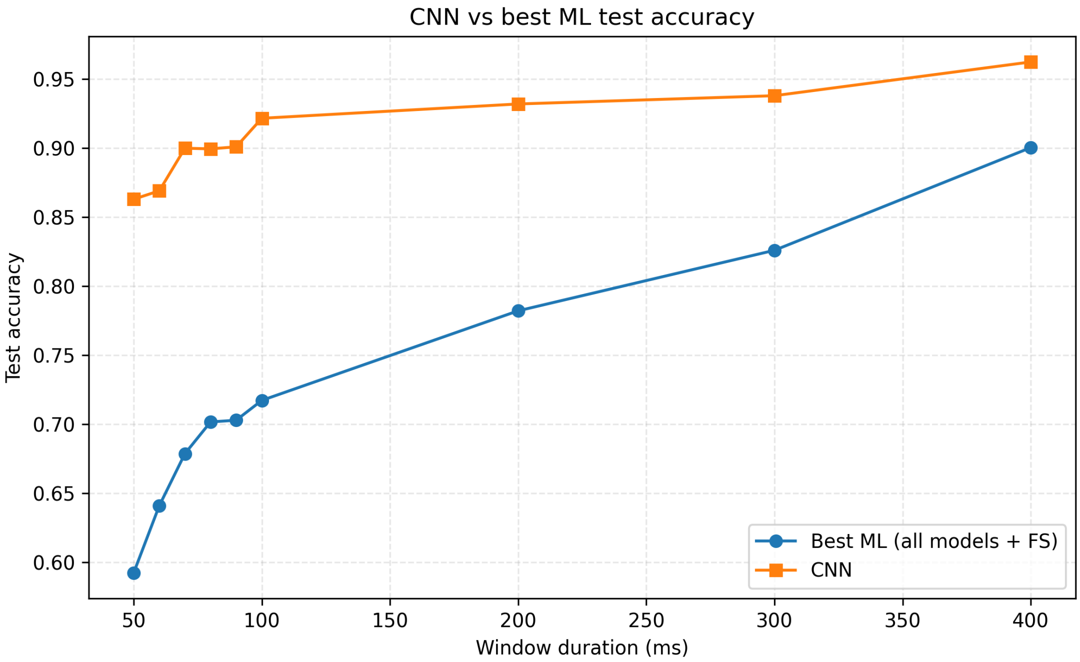

Figure 4.6: Comparison between the CNN and the best-performing ML configuration for each window duration T_w.

$\mathcal{F}_{T_w\text{-PCA}}$ and RF with $\mathcal{F}_{T_w\text{-RFE-RF}}$; for example, RF with RFE-RF features reaches 67.9% at 70 ms and 71.7% at 100 ms. For longer windows ($T_w \geq 200$ ms), RF combined with $\mathcal{F}_{T_w\text{-RFE-RF}}$ consistently provides the best ML performance, with test accuracies of 78.2%, 82.6%, and 90.0% at $T_w = 200$, 300, and 400 ms, respectively. Overall, across all windows and feature subsets, the feature-based approaches achieve test accuracies in the range 50.5%–90.0%, with the highest value obtained by RF with RFE-RF features at $T_w = 400$ ms.

In parallel, a convolutional model was trained directly on the windowed raw signals using the same train–validation–test split. Its test accuracy as a function of T_w is reported in Fig. 4.6, together with the best-performing ML configuration at each window. The CNN consistently outperforms the classical models for all window durations. At $T_w = 50$ ms, the CNN achieves a test accuracy of 86.3%, compared to 59.2% for the best ML configuration (KNN with $\mathcal{F}_{50\text{-PCA}}$), yielding an absolute gain of 27.1 percentage points. This advantage gradually decreases as more temporal information becomes available but remains substantial: at $T_w = 200$ ms the CNN reaches 93.2% versus 78.2%

Table 4.5: Summary of the statistical analyses performed for the 100 ms configuration in Chapter 4.

Analysis	**Test**	**Comparison**	***p*-value**	**Conclusion**
Classical pipeline comparison	Cochran's Q	SVM ($\mathcal{F}_{\text{100-RFE-SVM}}$), KNN ($\mathcal{F}_{\text{100-RFE-L1}}$), RF ($\mathcal{F}_{\text{100-RFE-R1}}$), SLFNN ($\mathcal{F}_{\text{100-RFE-L1}}$)	2.42×10^{-37}	Significant overall difference among the classical pipelines
Classical pipeline comparison (post-hoc)	McNemar + Holm	RF ($\mathcal{F}_{\text{100-RFE-RF}}$) vs SVM ($\mathcal{F}_{\text{100-RFE-SVM}}$)	1.01×10^{-36}	RF significantly outperformed SVM
Classical pipeline comparison (post-hoc)	McNemar + Holm	RF ($\mathcal{F}_{\text{100-RFE-RF}}$) vs KNN ($\mathcal{F}_{\text{100-RFE-L1}}$)	1.52×10^{-5}	RF significantly outperformed KNN
Classical pipeline comparison (post-hoc)	McNemar + Holm	RF ($\mathcal{F}_{\text{100-RFE-RF}}$) vs SLFNN ($\mathcal{F}_{\text{100-RFE-L1}}$)	2.26×10^{-16}	RF significantly outperformed SLFNN
Final offline comparison	McNemar	RF ($\mathcal{F}_{\text{100-RFE-RF}}$) vs 1D-CNN	2.15×10^{-66}	1D-CNN significantly outperformed the best classical pipeline
Condition analysis	Likelihood-ratio test	TCP main effect	2.47×10^{-1}	No significant TCP effect
Condition analysis	Likelihood-ratio test	GF main effect	6.29×10^{-1}	No significant GF effect
Condition analysis	Likelihood-ratio test	TCP × GF interaction	6.94×10^{-1}	No significant interaction effect

for RF with $\mathcal{F}_{\text{200-RFE-RF}}$, and at $T_w = 400$ ms it attains 96.2% compared to 90.0% for the best ML configuration (RF with $\mathcal{F}_{\text{400-RFE-RF}}$). Averaged across all tested window durations, the CNN improves test accuracy by approximately 18 percentage points over the best feature-based model.

To support the descriptive comparison above with formal inferential evidence, statistical analyses were carried out on the selected 100 ms configuration. The main results are summarized in Table 4.5.

First, classical model selection at 100 ms was evaluated at the level of complete classical pipelines. Since feature selection in this chapter was implemented in a classifier-specific manner, the inferential comparison focused on the final candidate pipelines rather than on a single pooled feature-selection factor. Cochran's Q test revealed a significant difference among the classical pipelines ($p = 2.42 \times 10^{-37}$). Post-hoc pairwise McNemar tests with Holm correction showed that all pairwise differences were significant, with RF trained on $\mathcal{F}_{\text{100-RFE-RF}}$ achieving the highest accuracy (71.73%), followed by KNN trained on $\mathcal{F}_{\text{100-RFE-L1}}$ (65.62%), SLFNN trained on $\mathcal{F}_{\text{100-RFE-L1}}$ (59.98%), and SVM trained on $\mathcal{F}_{\text{100-RFE-SVM}}$ (54.64%). The best classical pipeline was then compared against the 1D-CNN using McNemar's test. The result was highly significant ($p = 2.15 \times 10^{-66}$), confirming that the 1D-CNN (92.15%) significantly outperformed the best classical machine-learning pipeline (71.73%).

These results indicate that feature-based classifiers can exploit longer windows to achieve competitive performance, especially RF trained on the RFE-RF subset, yet the end-to-end CNN provides consistently higher accuracy at all temporal resolutions. For this reason, the CNN trained at the selected window duration of 100 ms was adopted as the predictive model in the real-time classification pipeline.

4.4.3 Scenario-wise performance across TCP and GF conditions

To further characterise the behaviour of the selected CNN at the window duration $T_w = 100$ ms, we analysed its performance as a function of the experimental scenario, defined by the combination of target closed position (TCP $\in \{1, 2, 3\}$) and grasping frequency (GF $\in \{0.2, 0.3, 0.4, 0.5\}$ Hz). For each TCP–GF pair, a confusion matrix was computed on the corresponding subset of the test data, considering the five hardness classes $H3, H5, H7, H10, H12$. The resulting 3×4 grid of confusion matrices is shown in Fig. 4.7, with entries reported as row-normalised percentages.

Across all scenarios, the CNN maintains a high level of performance, with scenario-wise test accuracies ranging from 83.5% to 96.1%. The lowest accuracy is obtained for the condition TCP 1, GF 0.2 Hz (83.5%), whereas the best-performing scenario is TCP 3, GF 0.3 Hz, where the accuracy reaches 96.1%. When aggregated over all TCP and GF conditions, the overall test accuracy at $T_w = 100$ ms is 92.1%, in line with the global result reported in the previous subsection.

To assess whether these descriptive variations corresponded to statistically meaningful condition effects, an additional trial-level analysis was performed on the final 1D-CNN. Since TCP and GF are defined for the original grasping trial, window-level predictions were first aggregated by averaging the confidence vectors of all windows belonging to the same trial, and the final trial label was assigned as the class with the

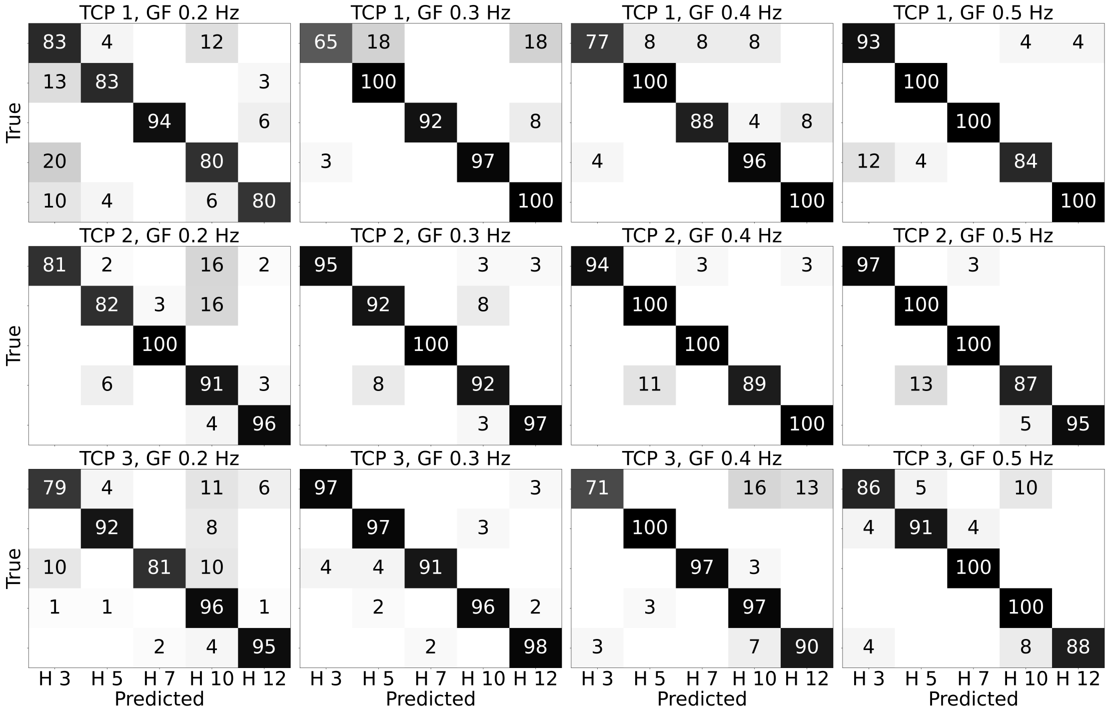

Figure 4.7: Scenario-wise confusion matrices for the CNN at window duration T_w = 100 ms. Rows correspond to target closed position (TCP levels 1–3) and columns to grasping frequency (GF ∈ {0.2 – –0.5} Hz)

highest mean confidence. The resulting trial-level accuracies remained uniformly high across all tested levels, reaching 98%, 100%, and 99% for TCP levels 1, 2, and 3, respectively, and 98.67%, 98.67%, 98.67%, and 100% for GF values of 0.2, 0.3, 0.4, and 0.5 Hz, respectively. However, formal likelihood-ratio tests did not reveal any significant effect of TCP ($p = 0.247$), GF ($p = 0.629$), or their interaction ($p = 0.694$) on classification correctness. Therefore, although some window-level scenario-wise fluctuations can be observed descriptively, the final CNN can be considered robust across the tested operating conditions at the trial level.

Analysing the per-class recalls averaged across all scenarios reveals a consistent descriptive pattern. The intermediate and higher hardness levels are classified very reliably, with mean recalls of 94.7% for $H5$, 95.2% for $H7$, 92.1% for $H10$, and 94.8% for $H12$. In contrast, the softest object $H3$ is more challenging, with an average recall of 84.7%. This tendency is particularly visible in some scenarios, such as TCP 1, GF 0.3 Hz and TCP 3, GF 0.4 Hz, where the recall for $H3$ decreases to 64.7% and 71.0%, respectively. In most cases, misclassifications of $H3$ are biased towards harder classes, especially $H10$ and $H12$, indicating a descriptive tendency of the model to overestimate hardness for the softest stimuli under some operating conditions.

A complementary descriptive pattern emerges from the confusion matrices of the hardest class, $H12$. Its recall remains consistently high across scenarios, often at or above 90%, with the majority of errors corresponding to confusions with $H10$. Large confusions between intermediate hardness levels (e.g., $H5$ vs. $H7$) are rare, and the corresponding matrix entries remain small across TCP–GF conditions. This suggests that the learned representation largely preserves the ordinal structure of the hardness levels, such that when errors occur, they mostly involve neighbouring classes rather than arbitrary swaps.

Overall, Fig. 4.7 shows that the CNN provides robust hardness discrimination across the full range of TCP and GF settings. While some descriptive degradation can be observed in specific scenarios, particularly for the softest class $H3$, the formal trial-level analysis indicates that these variations do not translate into statistically significant effects of TCP, GF, or their interaction. These results support the suitability of the selected CNN as the core of the real-time pipeline under varying contact dynamics.

4.4.4 Probabilistic analysis and confidence calibration

For the best-performing CNN at $T_w = 100$ ms, we further analysed the probability output of the softmax layer to assess how reliable the predicted confidence is and how errors are distributed across classes. Each test window produces a probability vector $\mathbf{p} \in \mathbb{R}^5$ over the hardness classes $H3, H5, H7, H10, H12$ (corresponding to 3%, 5%, 7%, 10%, and 12% infill), from which we define the maximum probability $p_{\max} = \max_c p_c$, the probability assigned to the true class p_{true}, and the rank of the true class in the ordered probability vector.

4.4.4.1 Calibration of Predicted Probabilities

To assess the reliability of the predicted probabilities, we computed a calibration (reliability) curve by binning the maximum predicted probability $p_{\max}$ into ten equally spaced confidence intervals and measuring the empirical accuracy in each bin (Fig. 4.8).

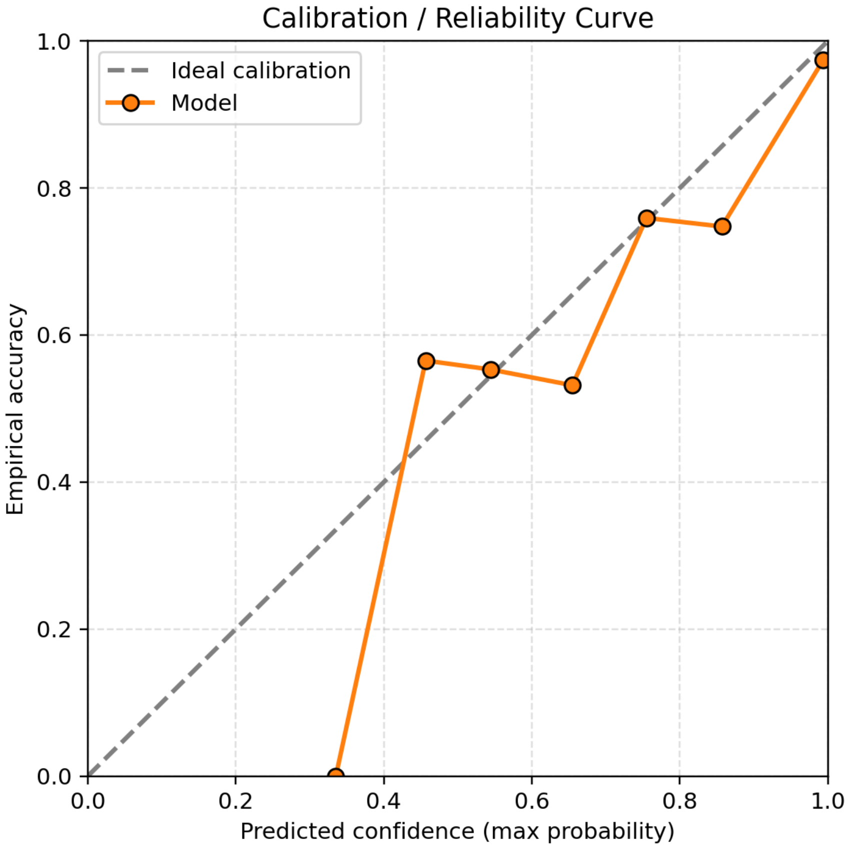

Figure 4.8: Reliability diagram for the CNN at $T_w = 100$ ms. Test predictions are grouped by maximum softmax probability, and, for each bin, the mean confidence is plotted against the empirical accuracy. The dashed line represents perfect calibration.

For low to intermediate confidence values ($p_{\max} \in [0.3, 0.6)$), the model exhibits moderate calibration variability: for example, predictions with mean confidence around 0.46 achieve an empirical accuracy of approximately 0.56, whereas those with mean confidence around 0.56 are correct only in about 0.38 of the cases.

In the higher confidence regime the behaviour stabilizes: predictions with mean confidence around 0.65 and 0.75 achieve empirical accuracies of roughly 0.70 and 0.66, respectively, while very confident outputs with $p_{\max} \in [0.9, 1.0]$ (mean confidence ≈ 0.99) are correct in 97.5% of the windows.

Overall, the CNN is reasonably well calibrated in the high-confidence region and slightly conservative in the mid-confidence range. From an application perspective,

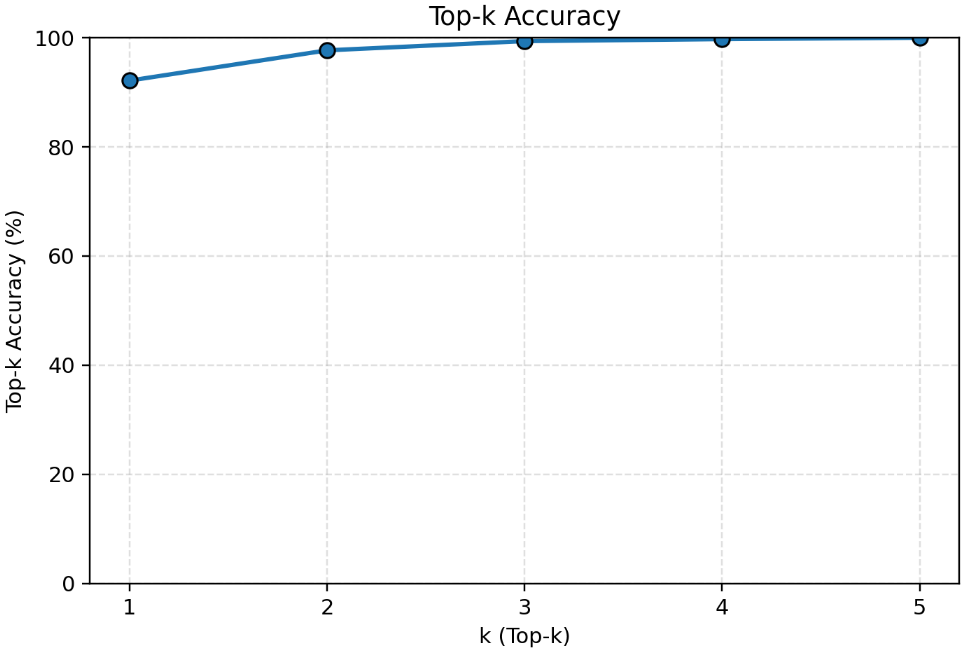

Figure 4.9: Top-k test accuracy of the CNN at $T_w = 100$ ms.

this suggests that a confidence threshold in the range $p_{\max} \geq 0.9$ can be used in the real-time pipeline to trigger highly reliable decisions, whereas predictions with lower confidence should either be combined across multiple windows (e.g., via majority voting) or discarded to avoid unstable classifications.

4.4.4.2 Top-k accuracy and near-miss behavior

The probabilistic behaviour of the CNN trained on 100 ms windows was first analysed in terms of Top-k accuracy (Fig. 4.9). On the test set ($N = 895$ windows), the standard Top-1 accuracy, which corresponds to the performance reported in the previous section, reaches 91.4%. Allowing the true class to appear within the top two predicted probabilities increases the accuracy to 96.8%, while Top-3, Top-4, and Top-5 accuracies reach 98.7%, 99.6%, and 100%, respectively. This indicates that when the network makes a mistake, the correct label is typically very close to the decision boundary and still receives a relatively high probability.

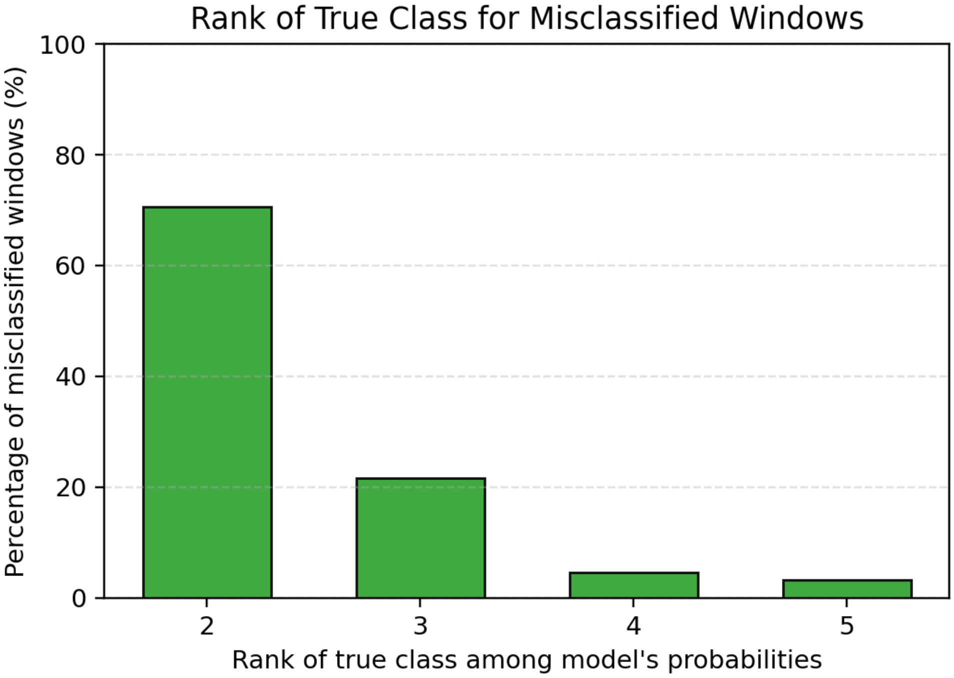

Figure 4.10: Rank of the true class among the model probabilities for misclassified test windows.

This observation is further quantified in Fig. 4.10, which reports the distribution of the rank of the true class among the model's probabilities for misclassified windows only ($N_{\text{mis}} = 77$). In 62.3% of these cases the true class is ranked second, in 22.1% it is ranked third, in 10.4% it is ranked fourth, and only in 5.2% of the misclassified windows the true class appears as the lowest-ranked label. Consequently, the true hardness class lies within the Top-3 predictions in 84.4% of misclassified windows and within the Top-4 in 94.8% of them. Taken together, these results show that most errors are "near misses", and suggest that Top-k or voting-based strategies can effectively exploit the full probability vector to increase robustness in the online setting.

4.4.4.3 True-class confidence distribution

Finally, we examined how the model distributes probability mass across the five hardness classes by inspecting the probability assigned to the true class p_{true} for each test window.

Figure 4.11: Mean confidence assigned to the true class, aggregated by hardness level $H1$–$H5$, with error bars indicating one standard deviation.

Figure 4.11 summarizes, for each hardness level $H1$–$H5$, the mean confidence assigned to the true label together with one standard deviation.

The mean confidence values are 82.7% for $H1$, 90.4% for $H2$, 91.4% for $H3$, 84.0% for $H4$, and 94.6% for $H5$, with relatively large standard deviations (between $\approx 22\%$ and $\approx 30\%$) reflecting the variability across indentation conditions. Despite this variability, the medians for all classes are close to or above 95%, indicating that when the network predicts the correct hardness, it usually does so with high confidence.

The combination of (i) high Top-k accuracy, (ii) a well-behaved calibration curve in the high-confidence regime, and (iii) consistently strong confidence on the true class across hardness levels supports the use of probability-based decision rules in the online pipeline. In particular, enforcing a minimum confidence threshold and exploiting the full probability vector (e.g., via Top-k aggregation across consecutive windows) appears to be a principled strategy to obtain reliable hardness discrimination in real time.

4.4.5 Real-time evaluation on the prosthetic hand

To assess whether the performance observed in the offline setting transfers to an online scenario, the CNN trained on $T_w = 100$ ms windows was deployed in the real-time pipeline and tested on the Hannes hand. Five hardness levels were considered, corresponding to the same 3%, 5%, 7%, 10%, and 12% infill objects used in the offline experiments (denoted $H3, H5, H7, H10, H12$). For each hardness level, three independent grasps were performed, yielding a total of 15 grasp events. During each grasp, the piezoelectric signals were streamed in real time, preprocessed as described in Section 4.3.5, segmented into fixed-size *non-overlapping* 100 ms windows, and fed to the CNN to produce a softmax probability vector $\mathbf{p} \in \mathbb{R}^5$.

Two complementary performance measures were extracted from the logged data: (i) an *event-level* confusion matrix, obtained by assigning a single hardness label to each grasp event via majority voting over all windows within that event; and (ii) a *confidence-gated* analysis based on the threshold motivated by the calibration study (Section 4.4.4.1). For each window, we define $p_{\max} = \max_c p_c$ and $\hat{c} = \arg\max_c p_c$. For the confidence-gated analysis, an event is counted as *detected* if it contains at least one *correct* high-confidence window, i.e., a window such that $\hat{c}$ matches the true hardness label and $p_{\max} \geq 0.9$. The detection latency is defined as the time elapsed between the start time of the first processed window and the end time of the first correct high-confidence window.

4.4.5.1 Event-level accuracy

Figure 4.12 reports the event-level confusion matrix obtained by majority voting over all windows within each grasp, without applying any probability threshold. Across the 15 grasp events, 12 were classified correctly, corresponding to an overall online accuracy of 80.0% at the event level. For three of the five hardness levels, all three grasps were correctly recognized (event-wise recall 100%), while the remaining two levels exhibited recalls of 66.7% and 33.3%, respectively. Importantly, the observed errors remain limited and do not include gross swaps between the softest and hardest objects.

Although the overall event-level accuracy (80.0%) is lower than the offline window-level Top-1 accuracy at $T_w = 100$ ms (91.4%, cf. Fig. 4.9), this difference is expected. First, the online experiment involves only 15 grasps (3 per hardness), so each error

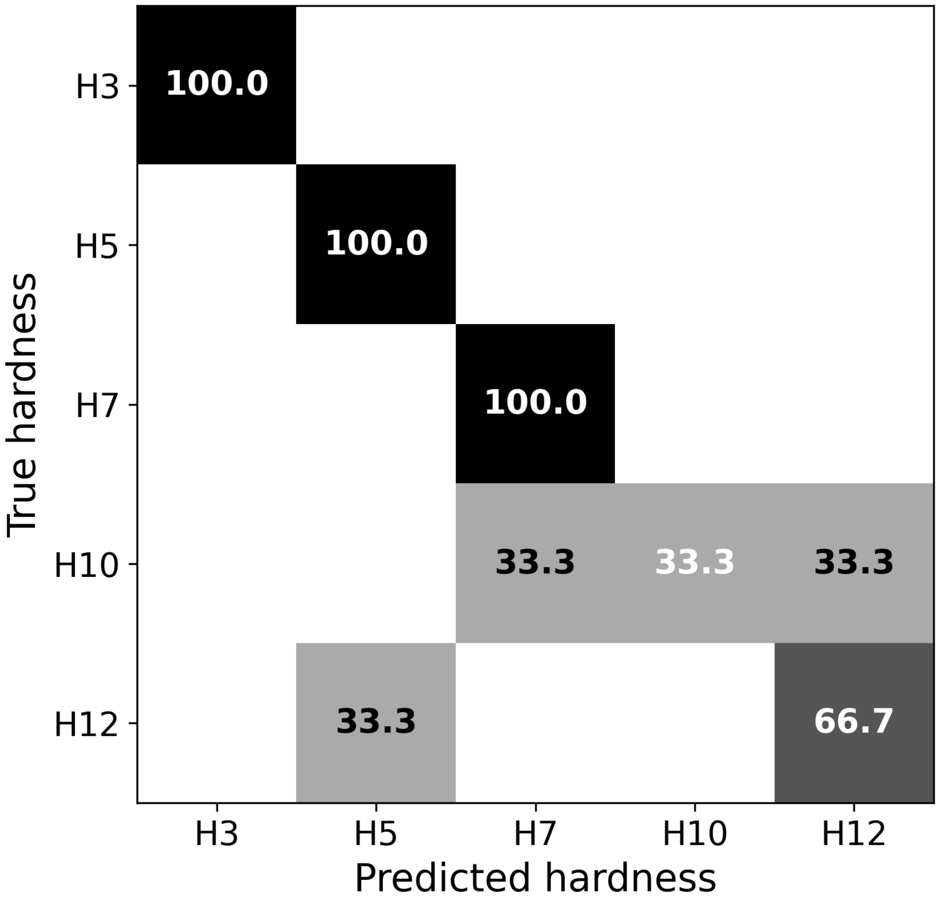

Figure 4.12: Event-level confusion matrix for the online experiment with the CNN at $T_w = 100$ ms. Each entry corresponds to one grasp event; rows indicate true hardness and columns predicted hardness. Entries are row-normalized and reported as percentages (three events per class).

changes the overall accuracy by more than 6 percentage points. Second, each event aggregates multiple windows; hence, a single misclassified grasp has a stronger impact at the event level than individual window errors in the offline setting. Finally, the online grasps were executed naturally, without enforcing the balanced scenario distribution used in the offline training and test sets, which may introduce additional variability in contact conditions.

Table 4.6: Online detection statistics for the CNN at $T_w = 100$ ms. An event is considered detected if it contains at least one *correct* window satisfying $p_{\max} \geq 0.9$. Latencies are measured from the start of the first processed window to the end of the first correct high-confidence window.

Class	N_{events}	N_{detected}	Detection rate (%)	Median latency (s)
$H3$	3	3	100.0	0.21
$H5$	3	3	100.0	0.11
$H7$	3	3	100.0	0.32
$H10$	3	1	33.3	1.91
$H12$	3	2	66.7	0.85
All	15	12	80.0	0.21

4.4.5.2 Confidence-gated detection and decision latency

To quantify how quickly reliable evidence can emerge online, we analysed the same 15 grasps under the confidence-gated criterion $p_{\max} \geq 0.9$, while requiring correctness ($\hat{c}$ equals the true label) as defined above. Table 4.6 summarises, for each hardness level and overall, the number of grasp events N_{events}, the number of detected events N_{detected}, the detection rate, and the median detection latency.

Across all hardness levels, 12 out of 15 grasps exhibited at least one correct high-confidence window, resulting in an overall detection rate of 80.0%. For the softer/intermediate objects ($H3$–$H7$), all grasps were detected with typical decision latencies below 0.35 s. For the two hardest objects ($H10$ and $H12$), detection rates decreased and latencies increased, indicating that these conditions may require longer temporal integration before the network produces a correct high-confidence output.

It is important to note that a confidence threshold alone does not guarantee correctness in the online logs: windows with $p_{\max} \geq 0.9$ can still correspond to an incorrect hardness label, particularly for the hardest objects where confusions with neighbouring classes occur. For this reason, the latency statistics in Table 4.6 are computed *retrospectively* using the known ground truth, and quantify the time to the first *correct* high-confidence evidence within each event. This behaviour is broadly consistent with the offline reliability analysis, where very high-confidence outputs ($p_{\max} \in [0.9, 1.0]$) were correct in 97.5% of the test windows (Section 4.4.4.1), while also highlighting that the hardest conditions remain the most challenging under fully online variability.

Taken together, the online results demonstrate that the CNN trained on 100 ms windows can be deployed in a real-time setting. The event-level accuracy without thresholding (80.0%) is broadly consistent with offline performance when accounting for the small sample size and less controlled conditions. Under the confidence-gated criterion, the system yields correct high-confidence evidence with sub-second latencies for the majority of grasps, supporting probability-based decision rules as a practical component toward robust hardness discrimination in real time.

4.5 Discussion

In the broader landscape of tactile perception, hardness (or compliance) inference has been studied extensively in robotics as part of building artificial touch systems that can support dexterous manipulation, often leveraging dense tactile sensors and learning-based mapping from contact signals to material properties [187, **?**, 188]; however, translating these ideas to prosthetic hands introduces additional constraints that are not always addressed in controlled robotic benchmarks, including wearable integration, signal robustness, and variability induced by natural grasps. Within this context, the novelty of this chapter is not the use of a CNN per se, but the demonstration of an end-to-end tactile perception capability built around (i) a prosthesis-compatible sensorization approach (a shielded, glove-integrated piezoelectric patch used on the index fingertip), (ii) a structured yet realistic dataset that explicitly includes multiple hand-closure levels and grasping frequencies, and (iii) an evaluation that couples offline accuracy with decision reliability and timing considerations required by real-time operation. The offline results show that, under substantial variability in contact dynamics (TCP and GF), hardness information remains consistently recoverable from short non-overlapping windows, and that end-to-end representation learning yields strong performance even at short temporal contexts where feature-based models struggle; this supports the practical claim that the tactile transients captured by the high-bandwidth piezoelectric sensing are informative enough to discriminate hardness without relying on longer, more stationary segments. Importantly, the probabilistic analyses (calibration, Top-k behaviour, and confidence on the true class) clarify how the model fails: most errors are near misses where the true hardness still receives high rank, suggesting that probability-vector aggregation can be exploited beyond Top-1 decisions—an insight that directly motivates majority voting and confidence-gated evidence accumulation in the online loop. In real-time, the event-level results further expose a key prosthetic-relevant phenomenon: while high confidence is strongly associated with correctness in aggregate, it is not a guarantee under online variability, especially for neighbouring hard classes, hence the chapter's deliberate choice to define detection latency as the time to the first correct high-confidence window; this definition is conservative but appropriate for prosthetic assistance scenarios where incorrect, high-confidence outputs could be harmful if used for autonomous control or user feedback. The resulting detection/latency patterns also provide an interpretable direction for future improvement: the hardest objects (H10/H12) require longer temporal integration and show reduced detection rates, implying that either richer sensing coverage (e.g., multi-finger patches), improved domain robustness (e.g., training with more diverse contact realizations), or decision strategies that incorporate ordinal structure (since errors concentrate on neighbours) could further stabilize performance. Finally, relative to prosthetic research that prioritizes sensory restoration to the user through neural interfaces [196], this chapter complements that line of work by showing that fingertip sensorization can be used to infer actionable object properties under realistic grasp variability, creating a foundation for downstream applications such as adaptive grasp modulation, automatic parameter tuning of grasp controllers, and context-aware sensory feedback policies where the system can decide what to convey and when to convey it based on confidence and latency constraints.

4.6 Conclusion

This chapter presented a prosthetic-hand tactile perception pipeline for hardness discrimination based on a glove-integrated piezoelectric sensing patch used on the index fingertip of the Hannes hand. A dataset was collected under systematic grasp variability by combining five object hardness levels with multiple target closed positions and grasping frequencies, and tactile signals were processed using non-overlapping windows to study the trade-off between temporal resolution and classification reliability. Across window durations, an end-to-end 1D-CNN consistently outperformed feature-based classical models, and the selected 100 ms configuration achieved strong offline performance while remaining compatible with real-time use. Probabilistic analyses (calibration, Top-k accuracy, and confidence statistics) provided a principled basis for probability-based decision rules, motivating majority voting and confidence-gated evidence accumulation. Finally, real-time tests on the prosthetic hand demonstrated event-level feasibility and quantified decision latency under streaming operation, highlighting both the practical utility of confidence-aware inference and the increased difficulty of the most challenging conditions under online variability. Taken together, the results support the use of wearable piezoelectric fingertip sensorization as a viable route to robust hardness perception on prosthetic hands, and establish the methodological groundwork for future applications that exploit hardness estimates for adaptive control strategies and/or feedback delivery policies constrained by reliability and timing.

CHAPTER 5

End-to-End Tactile Sensing and Electrotactile Feedback system

5.1 Introduction

Modern upper-limb prostheses have progressed substantially in mechanical design, actuation, and myoelectric control; however, they still provide limited or no somatosensory information to the user. This missing feedback is not a minor add-on: in the intact neuromuscular system, cutaneous and proprioceptive afferent signals continuously inform grasp regulation, contact localization, slip detection, and interaction timing, and they are tightly integrated with motor planning and correction. Consequently, the absence of tactile sensation in prosthetic hands is strongly associated with increased visual attention demands, slower and less reliable manipulation, reduced embodiment, and user dissatisfaction. Recent reviews emphasize that despite strong laboratory evidence supporting sensory substitution or restoration, *most commercial prostheses still do not incorporate sensory feedback*, largely due to challenges that span hardware integration, robustness, usability, calibration burden, and the lack of standardized outcome measures that reflect real-world performance [189, 190].

Addressing these limitations requires moving beyond isolated components (a sensor here, a stimulator there) toward *end-to-end* closed-loop solutions that can operate in real time and remain stable across subjects and sessions. At a system level, an effective feedback pipeline for upper-limb prosthetics must (i) acquire tactile information with adequate spatial coverage and bandwidth, (ii) transform raw signals into robust contact descriptors under variable grasp dynamics, (iii) encode the information into stimulation patterns that are learnable and interpretable, and (iv) deliver those patterns through a wearable interface that is comfortable, safe, and scalable in channel count. The literature contains several important steps in this direction. Invasive and semi-invasive approaches (e.g., peripheral nerve interfaces) can elicit more natural sensations and have demonstrated functional benefits, but they require surgery and careful long-term management [197, 190]. Non-invasive approaches based on vibrotactile, mechanotactile, skin-stretch, or electrotactile stimulation are attractive for broader translation because they can be deployed without surgical intervention, but they face constraints in information bandwidth, spatial selectivity, and day-to-day repeatability [122, 189].

Within non-invasive strategies, multi-site stimulation has emerged as a key enabler of richer percepts and improved performance compared to single-channel cues. For example, multi-channel haptic feedback has been shown to unlock higher dexterity in demanding manipulation tasks when vision is limited, illustrating the importance of distributing feedback over the limb rather than compressing it into a single scalar signal [198]. In parallel, recent work has explored how to convey higher-level object interaction descriptors (e.g., contact extent) using electrotactile patterns, supporting the general notion that spatially structured stimulation can communicate interaction-relevant information beyond simple event alerts [199]. A particularly relevant milestone for prosthetic-oriented *distributed* feedback is the full-hand electrotactile framework that combines an electronic skin with matrix electrodes to achieve high-bandwidth human–machine interfacing, demonstrating the feasibility of scaling stimulation and sensing toward hand-like distributions [200].

Despite these advances, two gaps remain prominent when considering practical upper-limb prosthetic deployment. First, many studies focus on a single feedback variable (e.g., grip force magnitude or slip), whereas everyday grasping requires interpreting a mixture of *where* contact occurs, *how* contact spreads across digits, and *how* contact evolves over time. Second, only a limited subset of works demonstrate an integrated real-time pipeline in which distributed tactile sensing on a prosthetic hand is mapped in a structured manner to distributed stimulation on the user, with an experimental evaluation that spans multiple perceptual tasks. This motivates the need for systems that are simultaneously (i) high-channel and distributed, (ii) real-time and robust, and (iii) assessed across representative perceptual dimensions that matter for manipulation (spatial localization, contact extent, grasp state, temporal dynamics, and material-related cues) [189, 122, 190].

In this chapter, we address these challenges by presenting an end-to-end sensory feedback system that translates tactile events detected by a 64-sensor piezoelectric e-skin integrated on the fingertips and palm of a prosthetic hand into spatially organized electrotactile sensations delivered through flexible matrix electrodes. A key design principle is the nearly one-to-one mapping between tactile sensors, stimulation channels, and electrode pads, which preserves topographic organization and enables intuitive spatial coding while maintaining real-time operation. We evaluate the resulting pipeline in healthy subjects using five experiments designed to probe complementary perceptual dimensions: Finger/palm Detection (FD), object Size Detection (SiD), Closed-hand Position (PD), speed of contact (SpD), and Hardness Detection (HD). By combining distributed sensing, embedded acquisition, real-time event detection, and distributed stimulation with a multi-task psychophysical assessment, the proposed framework aims to fill the translation gap between laboratory demonstrations of isolated feedback cues and practical, scalable closed-loop feedback for upper-limb prostheses.

This chapter is organized as follows. Section 5.2 describes the system hardware and integration (prosthetic sensorization, interface electronics, and stimulation interface) and details the real-time signal processing and mapping strategy. Section 5.3 presents the experimental setup, protocol, and the five evaluation tasks. The quantitative results are then reported and analyzed in section 5.4 using aggregated performance metrics and confusion-matrices structures to highlight both strengths and failure modes of the encoding. Finally, in section 5.5 we discuss the implications of recent state-of-the-art sensory feedback systems and outline directions for improving temporal and material-property discrimination while preserving strong spatial performance.

5.2 Materials and Methods

5.2.1 System description

The proposed system (figure 5.1) includes: i) piezoelectric sensing arrays (Tactile sensing system),ii) embedded electronics for signal acquisition,iii) Hannes prosthetic hand,iv) electrotacile stimulator,v) flexible surface electrodes and, vi) host PCs, first PC (2.3 GHz, 32GB), second PC (3.0 GHz, 32GB).The e-skin converts the mechanical contact into a set of electrical signals (one signal per sensor). When the prosthetic hand perform a grasp and a contact occurs between the sensors and the object the sensor signals are then sampled by the embedded electronics, processed and sent to one of

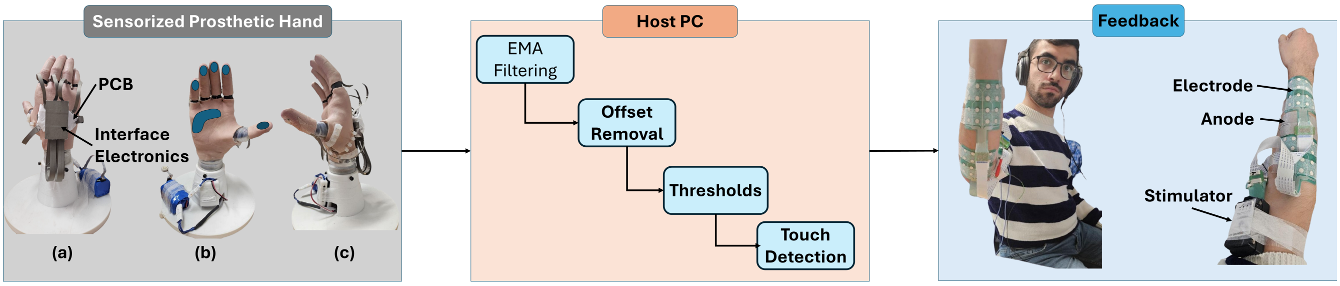

Figure 5.1: System Architecture. The system comprises an e-skin of 64 sensors, interface electronics for signal acquisition, and a multichannel stimulator with 4 flexible matrix electrodes integrating 64 electrode pads to deliver electrotactile stimulation to the subject. Consequently, the system translates the tactile data recorded by the e-skin into electrotactile stimulation delivered in real-time to the subject forearm.

the host PCs. The signals are then filtered and when sensor output exceeds a given threshold contact events are detected, and the activated sensors are highlighted on the graphical user interface on the PC (visual feedback). This information is used to set the state (on/off) of the corresponding stimulation channels (mapping between sensors and stimulation pads is presented in section5.2.3. The PC generates appropriate stimulation commands and sends them to the stimulator via Bluetooth. The electrotactile stimulation is delivered to the subject through matrix electrodes placed on the volar and dorsal sides of the hand. As the feedback pipeline runs in real-time, the tactile interaction recorded by the electronic skin is translated online into dynamic tactile sensations elicited across the subject's hand. 5.1 illustrates the online pipeline of the system.

5.2.1.1 Hannes hand Sensorization

1. **Sensing arrays:**

   The sensing technology used in chapters 3 and 4 was used in this activity. In this study, we used a fully screen-printed flexible sensor array based on piezoelectric polymer poly(vinylidene fluoride trifluoro-ethylene) (P(VDF-TrFE)) sensors, fabricated by JOANNEUM RESEARCH [201]. Each sensor is composed of three layers: a bottom electrode printed on a flexible PCB (100 μm thick), a ferroelectric polymer P(VDF-TrFE) layer (5.1 μm thick), and an electrode on top (PEDOT:PSS). The sensor is protected by a UV-curable lacquer layer. The whole structure was screen printed on a transparent and flexible PET (175 μm thick) sub-

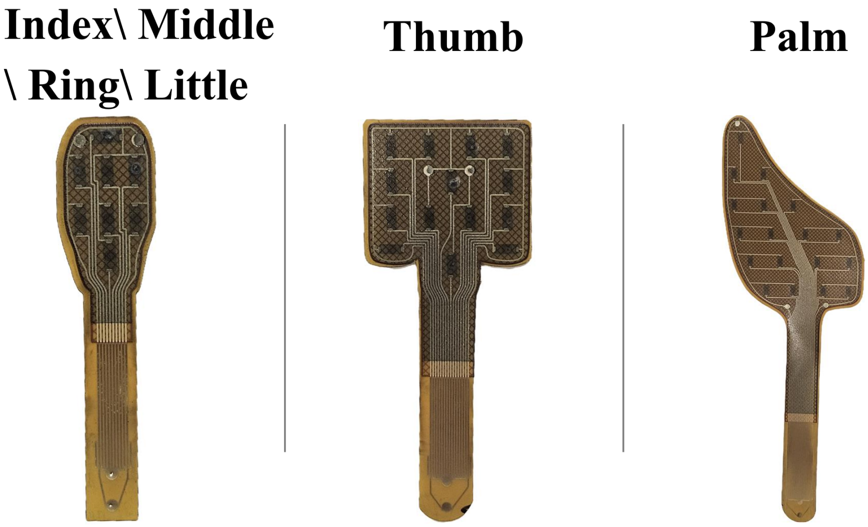

Figure 5.2: Image presents the geometry of the piezoelectric sensors used in this experiment. Also, it shows how the sensors are distributed in each sensing array. Moreover, it shows where each type of sensing array was integrated on the prosthetic hand.

strate. The fabrication technology was presented and validated in [41]. It is worth noting that piezoelectric polymer sensors have a bandwidth integrating those of all mechanoreceptors in the human skin. This kind of sensor can therefore measure dynamic contacts and infer information about static and dynamic contacts through data processing. A complete set of sensing arrays with different geometries and sensor distributions (range 8 - 18 sensors) was developed to fit the fingers and palm of the Hannes hand. Figure 5.2 shows the geometry, sensor distribution and size of the sensing arrays while figure 5.1.b highlights sensor distribution over the fingers and palm of the Hannes hand. The finger sensing arrays were designed to fit on the fingertips easily.

2. **Sensor Integration:**

   Following the activity developed in 4 and the developed sensorization in section 4.2.1.2, some challenges regarding the system were discovered. To extend the previous sensorization into a complete distributed tactile sensing system, additional sensing arrays were integrated on the hand, and further mechanical modifications were introduced in the glove structure. To facilitate mounting and removal of the glove on the prosthetic hand, the outer glove was mechanically modified. First, the finger portions of the outer glove were cut so that each finger became an independent sleeve rather than being connected to the main body of the glove. This modification simplified the positioning of each finger segment on the corresponding prosthetic finger. In addition, an opening was created along the radial side between the thumb and the wrist. Between the transparent inner glove and the outer glove, two 3D-printed silicone straps were integrated (figure 5.1 block Sensorized Prosthetic Hand.C), forming belt-like structures that wrap around the hand from the dorsal and lateral sides. Once the glove assembly was correctly positioned, these silicone "belts" were tightened and closed, thereby securing the outer glove to the inner glove and preventing relative motion between the glove and the prosthetic hand. Although the finger sleeves of the outer glove were no longer connected to the main glove body, the transparent inner glove fit tightly around the prosthetic fingers, which prevented the finger segments from being peeled off during grasping and object manipulation. After these mechanical modifications, the sensing arrays were integrated as described in Section 4.2.1.2. Each sensing array was glued to the transparent inner glove at the fingertip using silicone glue, and a second thin layer of the same glue was applied on top of the sensors before mounting the corresponding portion of the outer glove. As shown in Fig.5.3 a, 14 sensors were distributed over the thumb, while 8 sensors were placed on the fingertip of each of the index, middle, ring, and little fingers using identical sensing arrays. The remaining 18 sensors were positioned on the palm region to cover the palmar (volar) surface of the hand. A small PCB with four FPC sockets was fixed to the dorsal side of the hand inside a protective box using the same silicone adhesive tape and used to route the sensor signals to the interface electronics. The sensing arrays were connected to this PCB via flexible printed-circuit (FPC) tails, and hot glue was applied locally to stabilize the track–socket junctions mechanically.

5.2.1.2 Interface Electronics

The embedded electronics used in this study is an extended version of the systems presented in the previous chapters. Compared to the previous design that can accommodate up to 32 sensors, the current design can handle up to 64 sensors through two daisy-chained analogue-to-digital converters (DDC232, Texas Instruments) [202] mounted on the top and bottom side of the PCB. The two DDC232 were daisy-chained in order to avoid data loss during the acquisition process. Four sockets (two on the top side and two on the bottom side) acquire signals from 64 sensors where each block accommodates 16 sensors. The BL600 module is used to read, process and transmit sensor data. Our solution provides more channels with respect to analogous state-of-the-art electronic systems to acquire and process data from piezoelectric tactile sensors for touch sensing with robotic hands. In particular, the interface electronics for the system on chip device for prosthetic applications presented in [203] offers 13 channels only, while the voltage-mode approach to read PVDF taxel outputs proposed in [204] only manages data from nine sensors. In the present study, the embedded electronics was configured to collect and process tactile data from 64 sensors at 2K samples/s. The 2 kHz sampling rate was used to capture the full bandwidth of the sensor, which is beneficial for detecting the onset of contact events characterized by steep increases or decreases in the signal.

5.2.1.3 Feedback System

1. **Electrotactile stimulator:**

   The feedback interface employs a 64-channel programmable, battery-powered electrotactile stimulator (extended "Tactility", Tecnalia Research and Innovation, Spain [205]). This device is an upgraded version of the 32-channel stimulator used in previous work, developed by the same company and sharing the same electrical specifications [206]. The stimulator generates biphasic symmetric current pulses that can be distributed in time and space over 64 independently controllable channels. It is fully programmable, and the stimulation parameters can be adjusted online by sending text commands from the host PC via a Bluetooth connection. The amplitude of the current pulses can be modulated in the range 0–10 mA in steps of 0.1 mA; the pulse width can be set between 50 and 5000 μs in steps of 10 μs; and the stimulation frequency can be adjusted with 1 Hz resolution up to a maximum of 400 Hz. In the present work, 62 channels were configured as cathodes and 2 as anodes. Electrotactile feedback was delivered in a dynamic, event-based fashion: stimulation was enabled only when the sensor response crossed predefined thresholds associated with ongoing contact changes, and it was kept off when the sensor output remained within the non-active range. The PC sends stimulation commands to the stimulator to activate or deactivate the corresponding electrode pads according to these threshold-based events. The 64 output channels of the stimulator were connected to custom electrodes, as described in the next section.

2. **Electrodes:**

   In the present work, the stimulator was connected to four biocompatible flexible matrix electrodes produced by Tecnalia Serbia. Each electrode was fabricated using polyester substrate, an Ag/AgCl conductive layer, and an insulation coating covering the conductive leads. Each matrix integrated 16 stimulation pads arranged in a 4×4 grid, yielding a total of 64 pads across the four electrodes. The pads on the matrices were used as active stimulation sites to elicit electrotactile sensations on the skin of the user's hand, while two larger self-adhesive electrodes acted as common references (anodes) for the array. To improve electrode–skin contact and ensure stable current delivery, all stimulation pads were covered with conductive biocompatible hydrogel (AG725, Axelgaard, DK).

5.2.2 Signal Processing

A real-time algorithm was developed to extract binary contact information from the dynamic voltage signals produced by an array of piezoelectric sensors embedded in a prosthetic interface. Each sensing point within the array was treated as an independent input channel, enabling parallel processing and facilitating localized contact detection. The binary information, in this context, indicates the presence ("1") or absence ("0") of contact on the sensor surface.

The processing pipeline begins with an initialization phase wherein a fixed-length buffer of initial samples is collected for each channel. This buffer is used to estimate and remove the DC offset, calculated as the mean value per channel. The estimated mean is subtracted from all subsequently acquired data in that channel to eliminate baseline drift and inter-sensor bias. This step ensures that each channel operates on a zero-centered signal, allowing consistent interpretation of dynamic variations across different sensor locations.

Following offset correction, the data stream is filtered using an exponential moving average (EMA) filter. The EMA filter, selected because it is convenient for implementation, is expressed as [207]:

$$y[n] = \alpha x[n] + (1 - \alpha)y[n - 1], \tag{5.1}$$

where $x[n]$ is the current input, $y[n]$ is the current output, and $y[n - 1]$ is the previous output; α is a factor used to set the cut-off frequency. In this work, the parameter α was set to 0.09, which corresponds to a cut-off frequency of approximately 30 Hz. The EMA thus serves as a low-pass filter, attenuating high-frequency noise while preserving rapid transients associated with tactile events. This filtered signal is then used for further processing and decision-making.

To define detection thresholds, a buffer of size N samples (chosen by the user) is filled with data acquired in the absence of contact with the sensors, reflecting the non-contact state. In earlier implementations, the algorithm utilized the first and third quartiles (Q_1 and Q_3) to characterize this baseline; however, the updated approach adopts a simpler and more robust strategy by computing the minimum and maximum values within the baseline segment. These extrema are then scaled to derive asymmetric thresholds for event detection. Specifically, the upper threshold is defined as

$$T_1 = \alpha_{\text{th}} \cdot |\max|,$$

and the lower threshold as

$$T_2 = -\alpha_{\text{th}} \cdot |\min|,$$

where α_{th} is a scalar amplification factor (e.g., $\alpha_{\text{th}} = 7$). To prevent sensitivity to minor fluctuations, a deadband region is defined using broader limits based on scaled values of T_1 and T_2, ensuring that only significant deviations trigger event detection.

The thresholded signal is then processed using a finite state machine consisting of two primary states: "steady" and "active". While in the “steady” state, the system monitors the filtered signal and remains idle until the signal exceeds either T_1 or T_2. Once a threshold is crossed, the state transitions to "active", and a cumulative integration process begins. The cumulative signal $X(t)$ is computed iteratively as

$$X(t) = X(t - 1) + v(t), \tag{5.2}$$

where $v(t)$ is the filtered and thresholded signal at time t. This cumulative transformation enhances sustained signal deviations while suppressing high-frequency noise and transient oscillations. If the signal re-enters the deadband region, the state returns to “steady”, and a decay mechanism is applied to $X(t)$ using

$$X(t) = X(t) - \frac{X(t)}{\beta}, \tag{5.3}$$

where β is a decay factor (e.g., $\beta = 10$), gradually resetting the cumulative signal to zero to avoid long-term drift.

To detect binary events, a peak detection algorithm is applied to the absolute value of the cumulative signal. A running maximum is continuously updated during the "active" state. When the current signal magnitude falls below the running maximum, the algorithm interprets this as a potential local peak. To avoid redundant detections due to oscillations or noise, a refractory period is enforced using the parameter skipSamples, which defines the minimum number of samples to skip before allowing a new peak to be registered. This refractory mechanism is controlled via a decrementing counter skipCounter, which resets to skipSamples upon every confirmed peak and disables further detections until it reaches zero. As a result, a binary signal is created for each sensor, effectively representing the interaction occurring on its surface.

5.2.3 Mapping sensors to electrodes

In the present work, the number of stimulation channels of the extended electrotactile stimulator, the number of pads in the electrode arrays, and the number of tactile sensing elements on the prosthetic hand were approximately matched. This allowed us to implement a nearly one-to-one mapping between tactile sensor outputs, stimulator channels, and electrode pads, without the need to group multiple sensors or pads on the same channel. The cathodic outputs of the stimulator were routed to the pads of the four 4×4 electrode matrices, whereas an additional channel was connected to a large self-adhesive electrode used as a common anode to close the stimulation circuit.

Each tactile sensor in the glove was associated with a dedicated stimulation electrode and its corresponding cathodic channel. Sensors located on the thumb, fingers, and palm were mapped to electrode pads in an orderly fashion, preserving the relative organization

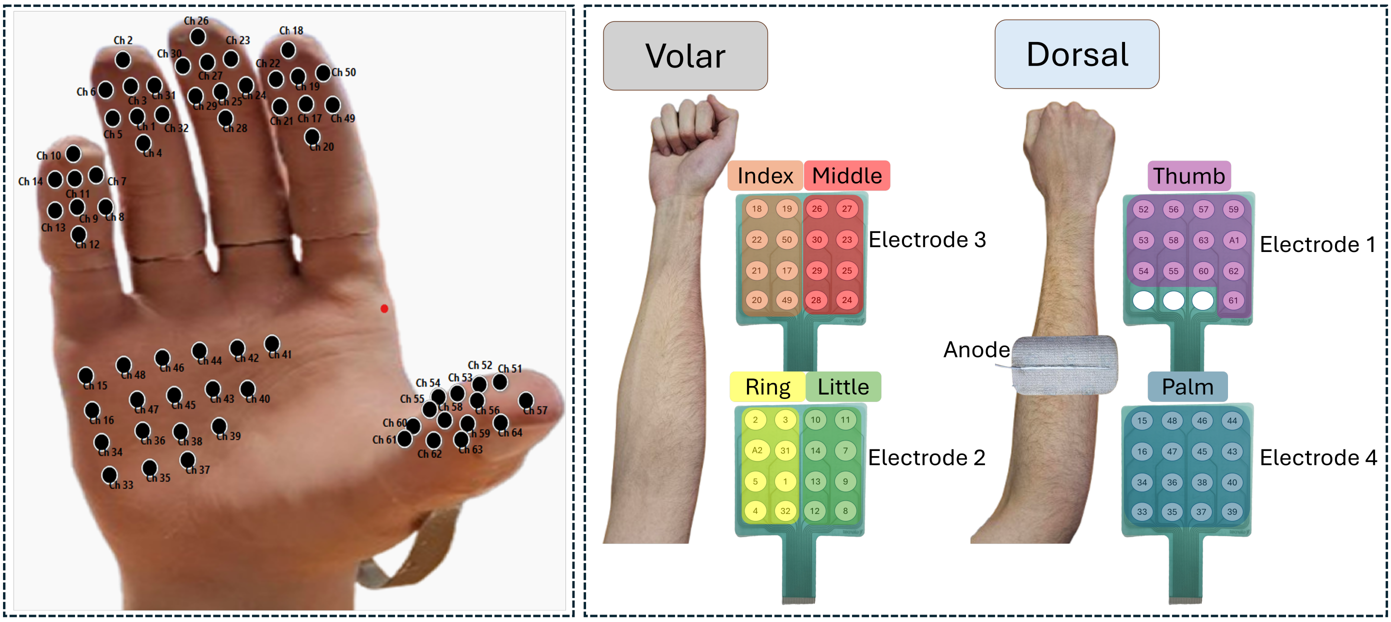

(a) (b)

Figure 5.3: (a) Sensor distribution within the sensorized glove of the Hannes hand. The sensors are numbered, and each sensor is associated to a stimulation pad on a matrix electrode on the subject forearm as shown in (b). (b) Electrodes and anode placement on the subject forearm, where 2 electrodes were placed on the dorsal side of the forearm and 2 on the volar side. The anode was placed on the dorsal side between the electrodes. Each stimulation pad has the number of the sensor that it represents.

of the digits and the palmar (volar) region (see Fig.5.3). For each sensor-channel-electrode triplet, the stimulation channel was activated whenever the corresponding sensor detected a contact event (binary state "1") according to the real-time algorithm described in Section 5.2.2, and was deactivated when the sensor returned to the non-contact state (binary state "0"). This one-to-one, event-based mapping ensured that local contact events on the prosthetic hand were encoded by spatially localized electrotactile sensations on the user's skin.

In addition to preserving a nearly one-to-one correspondence between sensors, channels, and pads, the spatial arrangement of the stimulation electrode on the forearm was defined according to a structured topographic transposition logic, rather than by arbitrary placement. The main objective was to preserve, as much as possible, the relative spatial organization of the tactile events detected on the prosthetic hand while keeping the resulting feedback easy to interpret and memorize. In particular, neighboring fingers on the hand were assigned to neighboring stimulation regions on the forearm, and the internal arrangement of the pads within each matrix was selected so as to preserve the orientation of the corresponding sensing elements. For example, when the matrix electrode was positioned on the volar side of the forearm, the stimulation points corresponding to the index finger were placed on the left side of the matrix, and those corresponding to the middle finger on the right side, so that the relative lateral order of the fingers was maintained after conceptually transposing the hand onto the forearm. Likewise, the stimulation points allocated to each finger were organized from top to bottom in a way that preserved the distal-to-proximal progression of the sensing elements, starting from the fingertip region and then extending toward the more proximal part of the finger. The thumb and palm regions were instead assigned to distinct areas, including the dorsal side of the forearm, in order to improve perceptual separability, avoid excessive crowding, and maintain a clearer spatial distinction for these functionally important hand regions. Therefore, the adopted layout is interpreted as a deliberate compromise between topographic fidelity, perceptual discriminability, and ease of memorization, rather than as the only possible mapping strategy.

Alternative mapping strategies are nevertheless possible and are important to consider when interpreting the present results. A first alternative is a regional grouping strategy, where multiple neighboring sensors are merged and drive a single pad or pad cluster; this reduces channel requirements and may increase robustness, but at the cost of spatial resolution. A second alternative is a hybrid spatial-intensity mapping, in which spatial position is preserved while stimulus amplitude, pulse width, or burst duration encodes an additional tactile variable such as contact strength or hardness. A third option is a sequential or scanning strategy, where a reduced number of pads are activated in time-multiplexed patterns to represent larger contact areas. The nearly one-to-one strategy was selected in this thesis because it offers the most direct way to test whether topographic information from a distributed e-skin can be preserved perceptually before introducing additional encoding dimensions or compression rules.

5.2.4 Rationale for the chosen forearm mapping

The adopted sensor-to-stimulation mapping was designed to preserve, as much as possible, the spatial localization of tactile events detected on the prosthetic hand while remaining easy for the user to learn and memorize. In particular, a nearly one-to-one mapping between sensing elements and stimulation pads was selected in order to convey not only the occurrence of contact, but also its approximate location on the artificial hand surface. The forearm distribution was therefore not chosen arbitrarily, but according to a structured topographic transposition logic: hand regions that are spatially or functionally related on the prosthesis were mapped to neighboring regions on the forearm, while maintaining a clear and interpretable organization of the stimulation layout. For this reason, the index and middle finger regions were positioned next to each other, and similarly the ring and little finger regions were kept adjacent, so that the relative neighborhood relationships of the fingers were preserved in the remapped feedback layout.

In addition, the chosen arrangement across the volar and dorsal sides of the forearm was motivated by perceptual clarity and channel separability. The thumb and palm regions were intentionally assigned to distinct forearm areas because they correspond to functionally important contact zones and include a sufficiently large number of sensing elements to justify a dedicated spatial representation. This separation was expected to facilitate interpretation of the evoked sensations and to reduce perceptual crowding, while remaining consistent with the one-to-one mapping philosophy adopted in this work. Thus, the selected forearm layout should be interpreted as a deliberate compromise between topographic fidelity, perceptual discriminability, and ease of memorization, rather than as the only possible mapping strategy. Other mapping organizations may also be feasible, and their comparison represents an important direction for future investigation.

5.3 Experimental Assessment

5.3.1 Experiments

Five experiments were designed to assess whether the feedback system can convey dynamic patterns in which contact changes over time (i.e., sliding across fingers, grasp characteristics). In one of the tests, the experimenter touched the e-skin, and in the rest,

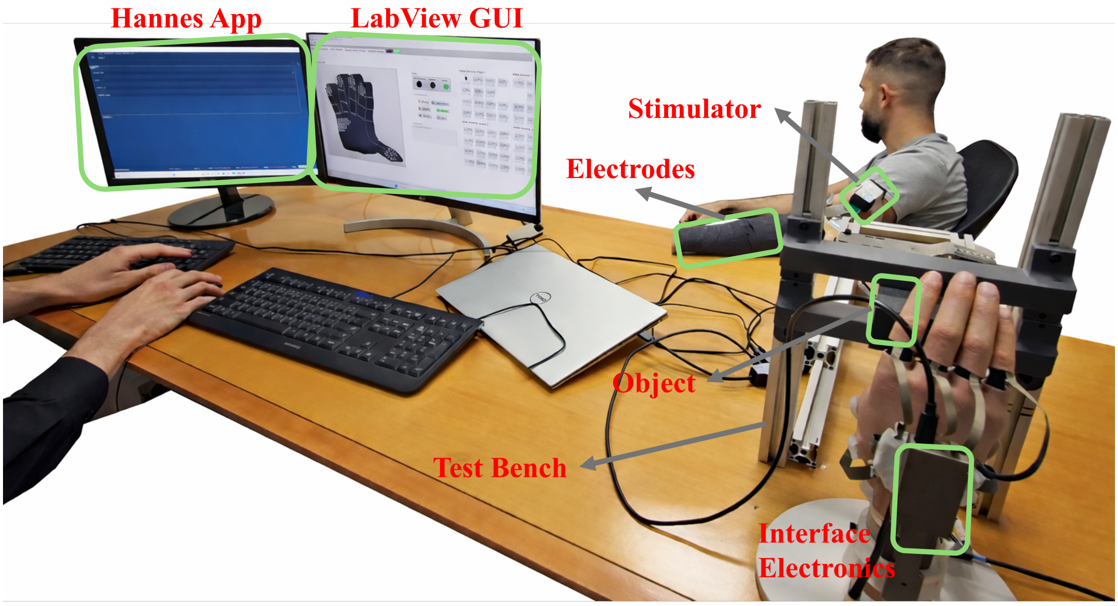

Figure 5.4: Experimental setup from experimenter perspective. Experimenter is interacting with the GUI's to control the prosthetic hand and monitor the performance of the end-to-end feedback system through the LabView GUI.

the hand was manipulated by the experimenter to grasp objects; those interactions with the sensors fired the detection algorithm, and touch was detected and transmitted to the subjects, who were asked to focus on the elicited sensations and interpret the feedback.

5.3.1.1 Finger Detection (FD):

This assessment aimed to evaluate whether the feedback system can successfully detect and convey to the subject the information on static contacts with low spatial resolution. The sensing areas of the e-skin patch were divided into 6 groups, where each group represented a sensing array. The 6 groups covered the volar side of the hand, starting by the fingertips of the fingers and the palm, as shown in figure 5.3.a. Similarly, 59 pads were selected on the electrodes and organized to form a mapping of the fingertips and palm in a close way to the hand, as shown in 5.3.b. Touch information was transmitted to the subject using spatial coding. Touch applied to one of the sensors on the e-skin was mapped into activating the corresponding stimulation pad in the matrix electrode.

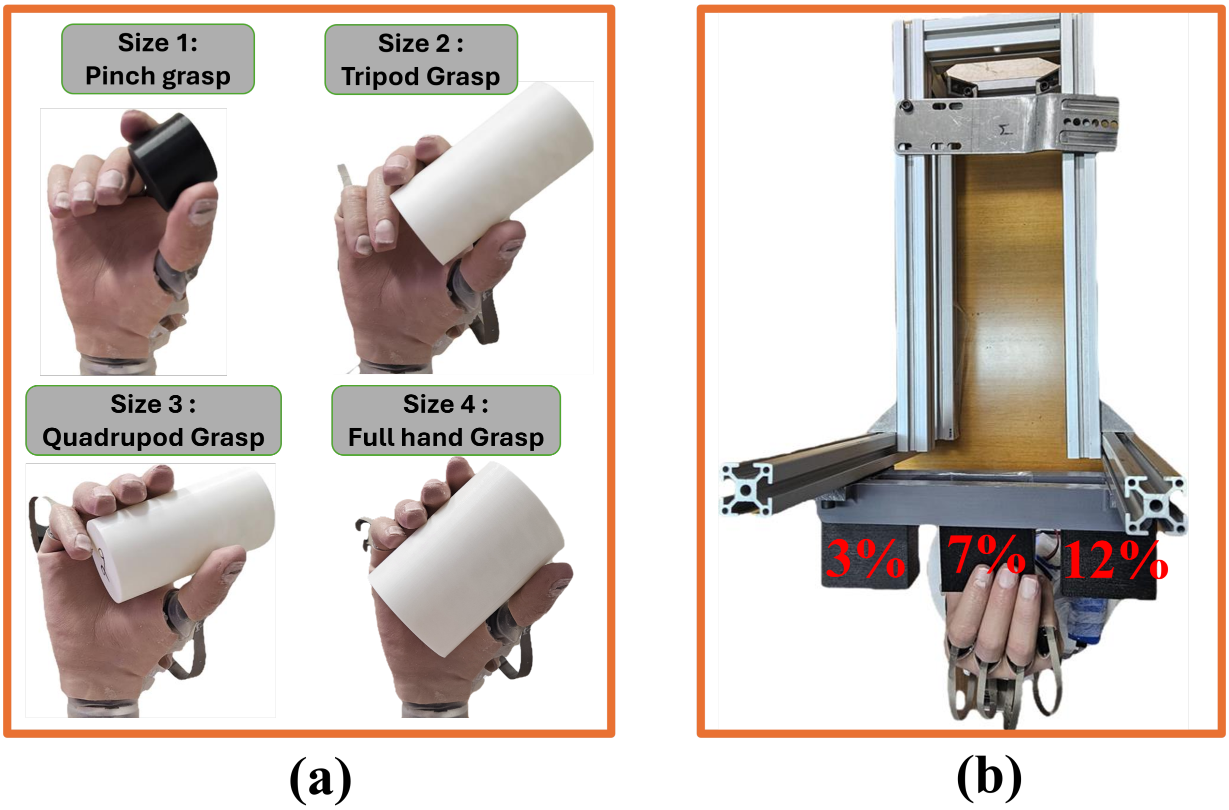

Figure 5.5: Experimental conditions. (a) Hannes hand grasping three objects, to manipulate 4 different levels of contact, which stimulates different combinations of sensors for each level. (b) Experimental conditions for the hardness detection experiment, where three objects of the same size and shape but different hardness are fixed on the same setup, and the hand moves horizontally in the same line between them.

5.3.1.2 Size Detection (SiD):

The goal of this experiment was to test the effectiveness of the feedback system in detecting and delivering touch information with higher resolution compared to that used in the FD experiment. In this case, the hand was used to grasp objects with different sizes (Figure 5.5 a). Grasping those sizes leads to activating stimulation spots on the electrodes (figure 5.3) depending on the activated sensors.

5.3.1.3 Closed Hand Position Detection (PD):

The goal of this experiment was to test the possibility of the feedback system in detecting and delivering touch information regarding the force applied when grasping. In this experiment, the hand was fixed to the table. A hard cube used in previous work [173] was fixed on the test bench as shown in figure 5.4 to ensure a consistent grasp. The hand was controlled to grasp this object with three different closed hand positions 70%, 85%, 100% (full closure) to ensure an additional squeeze on the object (more force), which increases the number of sensors in contact with the object, leading to activating more stimulation pads.

5.3.1.4 Speed of Contact Detection (SpD):

This experiment aimed to test the ability of the system to detect the speed of contact between the hand and the grasped object. The same setup condition used in the PD experiment was applied in this one, but instead of changing the closed hand position, the frequency of grasping was changed between 4 frequencies 0.02 (slowest), 0.05, 0.08, 0.1 (fastest) while fixing the closing hand position to 100%. This difference in the speed of contact will change how fast the stimulation is spread throughout the subject's hand, which helps him differentiate between the classes.

5.3.1.5 Hardness Detection (HD):

The goal of this experiment was to test the ability of the system to detect the hardness levels of the grasped objects. Three levels of hardness were tested using the 3D-printed cubes used previously [173] with the levels of filling 3% (Soft), 7%, 12% (Hard) (figure 5.5 b). The same setup used in PD and SD experiments was used in this one, but instead of fixing only one cube on the test bench, the three were fixed on the test bench, and the hand was moved by the experimenter in a parallel direction to the cubes and in similar positions for 3 of them.

5.3.2 Setup and Protocol

Five healthy subjects (male, $age 27 \pm 4$ years) participated in the five experiments described in subsection 5.3.1. Before starting, the subjects signed an informed consent form. Figure 5.4 shows the experimental setup used in all experiments. The experimenter was seated comfortably on a chair with their back towards the experimenter and the setup as shown in figure 5.4 so the subject couldn't see the hand nor the experimenter interacting with the hand. The forearm of the dominant hand was placed on the table, and the matrix electrodes were then positioned on the volar and dorsal side of the subjects forarm as shown in figure 5.3.b.

The hand was fixed on the table in front of the test bench holding the objects (figure5.4). The objects used to be held between the gray 3D-printed bars and fixed with double-sided tape. The interface electronics was connected to a host PC through a USB and paired with the stimulator through Bluetooth. Two screen monitors were used during the experiments, one to visualize the HannesApp and then help in controlling the prosthesis, and the other one to visualize the LabView GUI controlling the sensing system and the feedback system.

Prior to the experiments, the sensation threshold was determined for each of the 62 pads using the methods of limits by varying the pulse amplitude. The amplitudes were additionally fine-tuned by the experimenter until the subject reported that the perceived intensity was similar for all the pads. The pulse rate and pulse width were the same for all the channels and set to 30 Hz and 100 μs, respectively.

As explained before, the experiments aimed to assess the subject's ability to identify the contact patterns applied to the e-skin, captured by the integrated sensing system and delivered to the subject through electrotactile stimulation via electrode matrix. Each test started with an introductory phase, in which the subject was presented with an explanation of the working principles of the sensory feedback system and the feedback mapping. The same experimental protocol was followed in all the experiments, and it comprised familiarization, supervised learning, and a validation phase. In all phases, each touch pattern (class) was presented five times to the subject.

5.3.2.1 Phase 1: Familiarization

In the familiarization phase, the subjects received online visual feedback when seeing the activity done by the experimenter on the e-skin or the grasp applied by the hand, which shows the applied touch pattern (sensor activity) and the corresponding stimulation pattern (pad activity). The subject was asked to use the visual feedback to build a mental

Table 5.1: Validation-phase performance across subjects for the five discrimination tasks. Values are mean ± SD across $n = 5$ subjects.

Task	#Classes	Class definition (stimulus condition)	Chance (%)	Total trials	Accuracy (%)
FD	6	Thumb, Index, Middle, Ring, Little, Palm	16.7	138	94.7 ± 6.5
SiD	4	Size 1: Thumb+Index (pinch); Size 2: Thumb+Index+Middle; Size 3: Thumb+Index+Middle+Ring; Size 4: all fingers + palm contact	25	100	82.0 ± 5.7
PD	3	Closed hand position: 70%, 85%, 100% closure	33.3	75	82.7 ± 10.1
SpD	4	Grasp frequency: 0.02, 0.05, 0.08, 0.10 at 100% closure	25	100	71.0 ± 15.2
HD	3	Hardness: 3% infill (Soft), 7% (Medium), 12% (Hard)	33.3	75	76.0 ± 7.6

Note: FD includes 138 trials (one subject completed fewer valid validation trials for FD). All other tasks include the full planned number of trials.

mapping between the experienced sensation and the visual description (i.e. touched sensors, size of grasped objects).

5.3.2.2 Phase 2: Supervised training

In the supervised training phase, online visual feedback was removed, the contact patterns were randomly applied, and the subjects were asked to guess the applied patterns. The experimenter then provided verbal feedback on the correct answer. Specifically, the experimenter said “correct” if the subject successfully guessed the active pattern or “incorrect” and the experimenter provided the correct pattern to the subject verbally.

5.3.2.3 Phase 3: Validation

During this phase, the protocol was the same as during the supervised training, however, no feedback on the correct answer was given to the subject. The validation phase was the main part of the experiment and the results from this phase were used to assess the performance, while the previous two phases were used as the training.

5.4 Results

Performance was quantified in the validation phase (Phase 3), during which subjects received no feedback about the correct answer. Table 5.1 reports validation accuracy

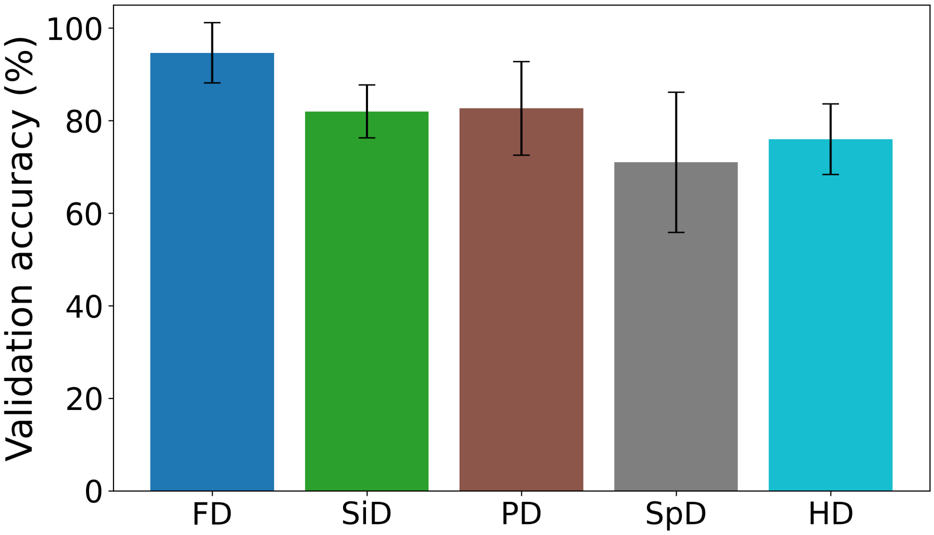

Figure 5.6: Mean validation accuracy across subjects for the five discrimination tasks. Error bars represent ±1 SD across subjects ($n = 5$).

(and balanced accuracy) across the five discrimination tasks, Fig. 5.6 visualizes the same task-level results, and Fig. 5.7 provides pooled row-normalized confusion matrices that reveal the dominant misclassification patterns. Since each class was presented five times per subject in the validation phase, the intended number of validation trials per subject was 30 for FD (6 classes), 20 for SiD and SpD (4 classes), and 15 for PD and HD (3 classes), resulting in the total trial counts shown in Table 5.1. Notably, FD includes 138 trials rather than 150 because one subject completed fewer valid validation trials for this task (see Table 5.1).

As shown in Table 5.1 and Fig. 5.6, FD achieved the highest performance (94.7 ± 6.5%), indicating that the proposed one-to-one sensor–pad mapping enables subjects to localize the stimulated anatomical group (Thumb, Index, Middle, Ring, Little, Palm) with high reliability, far above chance level (16.7%). The system also supported consistent discrimination of grasp-related spatial patterns in SiD (82.0 ± 5.7%) and PD (82.7 ± 10.1%), both well above their respective chance levels (25% and 33.3%). In contrast, tasks that relied more strongly on temporal cues or subtle material-property differences were more challenging, with SpD yielding the lowest mean accuracy (71.0±15.2%) and the largest inter-subject variability, and HD reaching 76.0 ± 7.6%. Balanced accuracy matched overall accuracy for all tasks (Table 5.1), confirming that the reported performance reflects consistent per-class recall rather than being driven by a single class.

The pooled FD confusion matrix in Fig. 5.7 exhibits a pronounced diagonal, with per-class recalls of 1.00 for Thumb and Middle, 0.957 for Ring and Palm, and 0.913 for Index and Little. Importantly, the few errors that occur are structured rather than random: Index is primarily confused with Middle (Index→Middle: 0.087), and Little is primarily confused with Ring (Little→Ring: 0.087), while Ring shows only minor confusion to Little (0.043) and Palm shows minor confusion to Thumb (0.043). This indicates that when subjects made errors, they tended to select a neighboring digit group, which is consistent with spatially adjacent stimulation patterns being perceptually more similar than non-neighboring ones.

A similar "neighbor-dominant" structure is observed for SiD in Fig. 5.7, where size classes correspond to progressively larger grasps: Size 1 (Thumb+Index pinch), Size 2 (Thumb+Index+Middle), Size 3 (Thumb+Index+Middle+Ring), and Size 4 (all fingers plus palm contact). The most distinctive condition is Size 4, which achieves the highest recall (0.92), consistent with the spatially extended activation pattern that includes palm stimulation. Conversely, the largest confusion occurs between Size 3 and Size 2 (Size 3→Size 2: 0.16), which is expected because these two conditions differ primarily by the additional involvement of a single digit (Ring). Additional errors remain concentrated among adjacent grasp sizes (e.g., Size 2→Size 1: 0.12 and Size 2→Size 3: 0.12), supporting that subjects interpreted SiD mainly through the spatial extent of the elicited stimulation.

For PD, which tested three closure levels (70%, 85%, and 100% closure), Fig. 5.7 shows that the full-closure condition is the most stable (Position 3 recall: 0.92), while the intermediate level is the most ambiguous (Position 2 recall: 0.72). This ambiguity manifests as confusions from Position 2 toward both extremes (Position 2→Position 1: 0.12; Position 2→Position 3: 0.16), while Position 1 is mainly confused with Position 2 (0.16). Therefore, PD errors are primarily boundary errors around the intermediate squeeze level rather than failures to detect closure-related stimulation changes.

Using the same setup but varying grasping frequency, SpD shows a clear trend in the confusion matrix (Fig. 5.7) whereby recall increases with speed: Speed 1: 0.64, Speed 2: 0.68, Speed 3: 0.72, and Speed 4: 0.80, indicating that faster grasping produced more distinctive temporal stimulation cues. Nevertheless, confusions remain

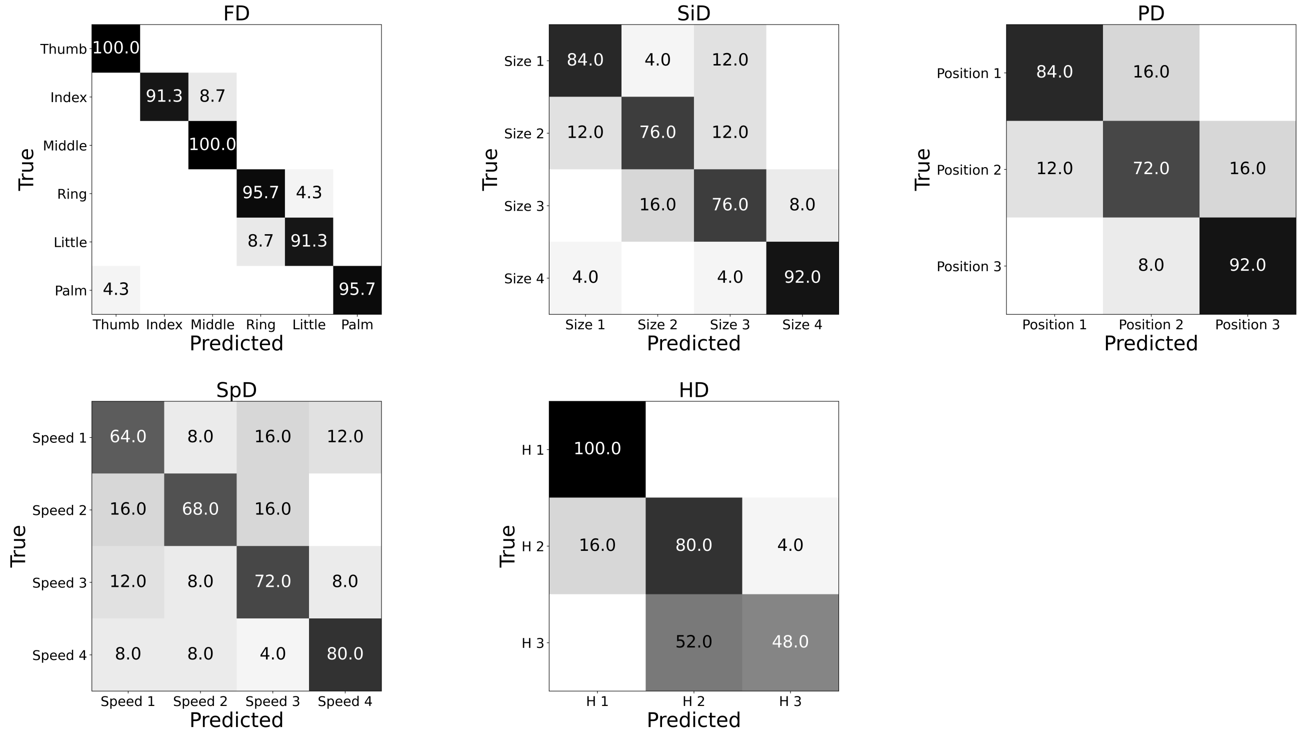

Figure 5.7: Confusion matrices pooled across subjects for FD, SiD, PD, SpD, and HD (validation phase). Each row sums to 1; diagonal entries indicate per-class recall.

concentrated among neighboring frequencies, especially for the mid-level condition (Speed 2→Speed 1: 0.16 and Speed 2→Speed 3: 0.16). Additionally, the slowest speed is occasionally confused with faster conditions (e.g., Speed 1→Speed 3: 0.16 and Speed 1→Speed 4: 0.12), which aligns with the larger variability observed at the task level (Table 5.1, Fig. 5.6).

Finally, HD exhibits an asymmetric confusion structure in Fig. 5.7 when discriminating 3% infill (Soft), 7% (Medium), and 12% (Hard). While Soft is always recognized correctly (recall: 1.00), Medium achieves recall 0.80 and Hard drops to recall 0.48, with a dominant error mode of Hard→Medium (0.52). This indicates that, under the current event-based encoding and grasp configuration, the perceptual cues for the two higher hardness levels overlap more strongly than those separating the soft level from the remaining classes.

5.5 Discussion

This work demonstrates an end-to-end, real-time sensory feedback pipeline that translates tactile events detected by a 64-sensor e-skin integrated on a prosthetic hand into spatially organized electrotactile sensations delivered through matrix electrodes using a nearly one-to-one sensor–channel–pad mapping. In the context of upper-limb prostheses, recent reviews emphasize that sensory feedback is critical for dexterous manipulation and embodiment, yet translation remains limited by the restricted information bandwidth of most interfaces and by the lack of standardized, functionally meaningful validation protocols [190, 189]. Against this backdrop, the quantitative results (Table 5.1, Fig. 5.6) and the structured confusion patterns (Fig. 5.7) show that the proposed encoding is particularly effective when the discriminative cue is primarily *spatial* (which anatomical region is active, or how spatially extended the contact is), and it becomes more challenging when the cue depends mainly on *temporal dynamics* (speed) or on subtle differences in interaction mechanics (medium vs. hard cubes).

The strongest outcome is the high FD performance (94.7 ± 6.5%), which provides direct evidence that the one-to-one mapping preserves a usable spatial organization in perception. Importantly, the FD confusion matrix reveals that errors are not arbitrary: the dominant confusions occur between anatomically neighboring digit groups (Index→Middle and Little→Ring), while long-range confusions across the hand are essentially absent (Fig. 5.7). This signature is consistent with a topographically coherent spatial code, where uncertainty tends to remain local rather than producing anatomically implausible mislocalizations. Notably, high-density distributed electrotactile feedback systems have previously shown that users can recognize static and dynamic contact patterns when sensing and stimulation are both spatially distributed [149]. The present results extend this evidence to a prosthetic-hand scenario and to a broader set of grasp-relevant discriminations (beyond contact location/sliding), supporting the idea that increasing the number of independently addressable channels can yield perceptually meaningful information—an aspect that is increasingly viewed as important for functional utility of feedback interfaces [189, 198].

In the same vein, SiD and PD both achieve ~82% mean accuracy (Table 5.1), and their confusion matrices are dominated by diagonal and near-diagonal terms (Fig. 5.7), indicating that subjects relied on ordered spatial cues. In SiD, the class definitions correspond to the number of digits involved in the grasp and the presence of palm contact; therefore, the elicited pattern should naturally expand as more sensors activate. The high recall for Size 4 (0.92) is consistent with this mechanism because palm involvement produces a spatially extended and highly distinctive percept, whereas the largest confusions occur between Size 2 and Size 3 (Fig. 5.7), which differ mainly by the additional involvement of a single finger. This observation aligns with prior psychophysical evidence that electrotactile interfaces can convey *contact size* effectively when the stimulus spatial extent increases in an orderly manner [208]. Similarly for PD, full closure (100%) is the most reliably identified condition (recall: 0.92), while the intermediate closure (85%) is most often confused with its neighbors (Position 2→Position 1/3), which is expected when the decision boundary depends on a gradual change in contact extent that can overlap due to small trial-to-trial variations in alignment and squeeze.

In contrast, SpD exhibits the lowest mean accuracy and the largest inter-subject variability (71.0 ± 15.2%, Table 5.1), which is coherent with the fact that speed discrimination is encoded predominantly through the *temporal evolution* of stimulation rather than through a static spatial footprint. The confusion matrix supports this interpretation by showing that the fastest condition is the most separable (Speed 4 recall: 0.80), whereas the intermediate speeds are most frequently confused with adjacent levels (e.g., Speed 2→Speed 1 and Speed 2→Speed 3) (Fig. 5.7). This pattern suggests that temporal cues become more salient as stimulation transitions become faster, while slower patterns are more sensitive to perceptual timing strategies and to minor variability in interaction dynamics. Finally, HD achieves moderate performance (76.0 ± 7.6%) but is limited by a strong overlap between the two higher hardness levels, as indicated by the dominant Hard→Medium confusion (0.52) and the reduced recall for Hard (0.48) (Fig. 5.7). Since the current pipeline encodes contact through binary events and spatial distribution, the perceptual differences between 7% and 12% infill cubes may be insufficient under the tested grasp configuration if both conditions generate similar contact footprints and event patterns at full closure. In this respect, the results suggest a clear direction for improvement that remains compatible with real-time constraints: augmenting the binary event representation with additional low-cost descriptors (e.g., event rate or cumulative activity) and/or introducing an additional stimulation coding dimension (e.g., modulating intensity with a hardness-related metric). Such multidimensional coding is consistent with broader trends in bidirectional prosthetic systems, where adding richer feedback variables (rather than purely spatial on/off cues) is often necessary to discriminate interaction properties beyond contact location [209, 189].

Overall, the main achievement of this study is that a distributed tactile interface coupled with a nearly one-to-one multichannel electrotactile mapping supports highly reliable finger/palm localization and enables meaningful discrimination of grasp patterns that modulate contact extent, while also revealing specific and interpretable bottlenecks for temporally driven (SpD) and material-property-driven (HD) discrimination that can be addressed in future iterations.

5.6 Conclusion

This chapter presented and experimentally validated an end-to-end, real-time sensory feedback system for upper-limb prosthetics that converts distributed tactile events detected by a piezoelectric e-skin integrated on a prosthetic hand into spatially organized electrotactile stimulation delivered through matrix electrodes. The proposed framework integrates (i) a distributed sensing layer (64 sensing elements covering the fingertips and palm), (ii) embedded acquisition and real-time processing to extract robust binary contact events from dynamic piezoelectric signals, and (iii) a nearly one-to-one mapping between sensors, stimulation channels, and electrode pads to preserve the topographic structure of contact. This design choice enables an intuitive spatial code, while maintaining real-time operation suitable for online interaction.

The experimental assessment with healthy subjects demonstrated that the proposed encoding reliably communicates multiple grasp-relevant perceptual dimensions. In the validation phase, subjects achieved high performance in finger/palm localization (FD), reaching 94.7 ± 6.5% accuracy across six classes, which confirms that the mapping preserves a coherent spatial organization that users can interpret with minimal ambiguity. Beyond pure localization, the system supported meaningful discrimination of contact extent and grasp state: object size/contact extent discrimination (SiD) and closed-hand position discrimination (PD) both achieved $\sim$ 82% mean accuracy (Table 5.1, Fig. 5.6), with confusion patterns dominated by diagonal and near-diagonal terms (Fig. 5.7), indicating that errors mainly occurred between neighboring levels rather than across distant classes. These outcomes provide evidence that distributed electrotactile patterns driven by distributed tactile sensing can convey not only *where* contact occurs, but also *how* contact spreads across the hand during grasping.

In contrast, the tasks that depended primarily on temporal evolution of stimulation (SpD) or on subtle differences in interaction mechanics (HD) were more challenging. Speed of contact discrimination (SpD) yielded 71.0 ± 15.2% accuracy and exhibited the largest inter-subject variability, consistent with the fact that speed is encoded mainly through perceived timing rather than static spatial footprints. Hardness discrimination (HD) reached 76.0 ± 7.6% accuracy, but the confusion structure showed a pronounced overlap between the two higher hardness levels, suggesting that binary event-based spatial coding provides strong separability for large differences (e.g., soft vs. non-soft) but may be insufficient to separate adjacent material conditions under the tested grasp configuration. Importantly, these limitations are not merely negative findings; they provide clear guidance for future refinement while preserving real-time constraints. Specifically, enriching the event representation with additional low-cost descriptors (e.g., event rate, cumulative activity, or simple amplitude-related features) and/or adding an additional stimulation coding dimension (e.g., intensity modulation) are promising strategies to increase separability for temporally driven and material-property-driven tasks without compromising the robust spatial performance demonstrated for FD, SiD, and PD.

Overall, the main contribution of this chapter is the demonstration of a scalable, distributed, and real-time end-to-end feedback pipeline that delivers multi-attribute tactile information through a non-invasive interface. By combining distributed piezoelectric sensing, real-time contact-event extraction, and a structured multichannel electrotactile mapping, the proposed system advances toward practical bidirectional prosthetic solutions that can convey multiple interaction-relevant cues rather than a single scalar feedback variable. Future work will focus on extending validation to amputee participants, assessing robustness under longer-term use and more ecologically valid manipulation tasks, and refining encoding strategies to convey temporal and material-related information better while retaining the strong spatial interpretability enabled by the one-to-one mapping.

CHAPTER 6

General conclusion

6.1 Conclusions

Upper-limb prosthetic hands have reached a high level of mechanical sophistication, yet their practical utility and long-term acceptance remain strongly limited by the absence of reliable somatosensory information and intuitive sensory feedback [9, 10, 210]. In the intact human sensorimotor loop, tactile afferents provide rapid, information-rich cues that support stable manipulation, grip-force regulation, and the perception of object properties [6, 5]. Motivated by this gap, this thesis investigated how a piezoelectric, PVDF-based tactile sensing approach can be progressively developed from a localized biomimetic sensing unit into a prosthetic-hand-level sensing framework, and how tactile data can be processed (via ML/DL and signal processing) to generate *feedback-ready* outputs that are compatible with real-time operation.

Across the experimental activities presented in Chapters 3–5, the main outcome of this thesis is the validation of an end-to-end tactile sensing pipeline that links: (i) practical sensorization of prosthetic-hand embodiments using PVDF tactile sensors and dedicated interface electronics [26, 27], (ii) online-capable tactile data processing and learning-based decoding to infer meaningful object-related variables (with emphasis on hardness), and (iii) a feedback-oriented interpretation of distributed tactile activity suitable for non-invasive stimulation strategies [13, 14, 16, 10]. Rather than treating tactile sensing, decoding, and feedback as isolated blocks, the thesis emphasizes their *co-design*: the characteristics of PVDF sensing and interface electronics directly shape the signal conditioning and event detection strategies, which in turn constrain what can be inferred online, and ultimately determine what tactile information can be encoded for feedback.

From biomimetic fingertip sensing to robust hardness discrimination (Chapter 3). The first part of the thesis established that PVDF-based tactile signals acquired during indentation contain sufficient discriminative information to recognize multiple levels of hardness under variability in indentation speed and load. Using statistical feature extraction and lightweight learning models (including an SLFNN), the study achieved high multi-class hardness discrimination performance while preserving a pipeline compatible with incremental operation. Beyond reporting aggregate accuracy, the analysis explicitly studied how performance evolves as the interaction unfolds by using incremental windowing strategies. Two complementary views were provided: (i) percentage-based segmentation, reflecting progress as a fraction of the trial duration, and (ii) fixed-step cumulative segmentation, reflecting progress in discrete observation increments. This incremental analysis demonstrates a practically important property for real-time prosthetic use: reliable hardness information can emerge early in the interaction and stabilizes as more tactile evidence accumulates. Finally, probabilistic analyses (Top-k accuracy, calibration, and rank structure of errors) showed that misclassifications tend to be near misses and that the classifier's confidence contains useful information for uncertainty-aware decisions. This supports the broader thesis claim that tactile decoding for prosthetics should not only provide point predictions, but also provide *interpretable confidence* that can guide downstream control or feedback policies [14, 9].

Hand-level sensorization and online ML/DL decoding in a real-time pipeline (Chapter 4). Building on the initial validation at the fingertip level, the second part of the thesis moved toward prosthetic-hand embodiments by integrating tactile sensing on the Hannes prosthetic platform, where three fingertip sensors were mounted (thumb, index, and middle), and the index sensor was used for the hardness-discrimination experiments. A key contribution of this activity was the construction of an online processing pipeline on a PC that combined data acquisition, filtering, drift handling, activity/contact detection, incremental windowing, and 1D-CNN-based hardness inference. This step is crucial for translation: many tactile classification studies demonstrate performance offline, whereas practical prosthetic usage requires stable real-time acquisition and decision-making under timing variability and signal nonstationarities [18, 28, 30]. By demonstrating that tactile learning models can be executed in a live loop with consistent pre-processing and event-triggered inference, this chapter validated the feasibility of moving beyond retrospective analysis toward online operation.

Distributed tactile sensing with feedback-ready outputs (Chapter 5). The final experimental activity addressed a core limitation of conventional sensory substitution systems: the restricted bandwidth of tactile information that can be extracted and transmitted when only a few sensing and stimulation channels are available [16, 10, 210]. This chapter presented a distributed sensorization of the prosthetic hand with a dense network of tactile sensors (64 sensing units distributed across fingertips and palm) and an equally dense stimulation concept, where each tactile sensor is mapped to a corresponding stimulation site. In contrast to hardness classification (Chapters 3–4), the pipeline here relied primarily on signal processing to extract spatial contact/activation patterns robustly in real time, and to translate these patterns into feedback-ready signals suitable for electrotactile stimulation [13, 14]. The system was validated in controlled laboratory experiments with healthy participants, demonstrating the feasibility of transmitting spatially distributed tactile events through a multichannel non-invasive interface. While the experimental context differs from at-home longitudinal studies, it provides an essential intermediate step: evaluating dense sensing and stimulation under realistic integration constraints (hand geometry, wiring, mounting robustness) and during physical interactions that generate the tactile distributions the pipeline must handle.

Overall contributions and positioning. Taken together, the thesis demonstrates a coherent progression: *(1) validate tactile signatures and decoding for hardness under variability; (2) implement online learning-based inference on a prosthetic-hand embodiment; (3) extend sensorization to a distributed hand-level system and produce feedback-ready outputs from dense tactile activity.* This progression directly addresses central requirements identified in the literature for practical sensory-enabled prostheses, including robustness under realistic variability, real-time feasibility, and feedback relevance [9, 10, 210]. In summary, the work contributes evidence that PVDF-based tactile sensing, combined with appropriate interface electronics and decoding strategies, can support both (i) object-property inference (hardness) and (ii) distributed contact encoding for sensory feedback, thereby paving the way toward higher-bandwidth and more useful tactile human–prosthesis interfaces.

6.2 Future Work

Although this thesis demonstrates the feasibility of the proposed tactile sensing and processing framework, several research directions remain open to strengthen its robustness, increase its practical relevance for end users, and expand its capabilities.

Extending robustness and coverage of prosthetic-hand sensorization. Future work should further improve the durability, maintainability, and scalability of the distributed sensorization. This includes investigating protective encapsulation strategies that preserve electromechanical coupling for PVDF sensors while improving resistance to wear, humidity, and day-to-day mechanical stress, as well as developing modular integration and calibration procedures that facilitate replacement and maintenance without degrading signal quality [18, 19, 25]. In addition, long-term stability studies (drift, aging, and repeatability) are needed to quantify performance degradation over time and to design compensation strategies within the acquisition and processing pipeline.

From offline decoding to generalizable online inference under real-world variability. The online pipeline demonstrated in this thesis can be strengthened by systematically studying generalization across wider conditions: different object sets, contact locations, grasp types, and user interaction styles. Methodologically, this opens directions such as domain adaptation, confidence-aware decision thresholds, and continual recalibration of pre-processing parameters (e.g., drift baselines) to stabilize real-time inference. For hardness discrimination specifically, future studies can explore multi-task learning and sensor fusion (e.g., combining tactile data with prosthesis proprioception, motor current, or kinematic features) to reduce ambiguity among intermediate stiffness classes and to improve robustness to uncontrolled interaction profiles [28, 30].

Closing the loop: tactile-informed prosthesis control. A central next step is to use tactile inference not only for recognition or feedback, but also to enhance prosthesis *control*. The outputs already demonstrated in this thesis naturally support several levels of control integration:

> **Reflex-level control:** fast contact/slip-related indicators extracted from tactile signals can trigger rapid grip-force adjustments to stabilize grasp (anti-slip reflex) without requiring user interpretation.
>
> **Supervisory control using hardness inference:** predicted hardness (and associated confidence) can adapt grasp parameters such as closing speed, force ramp rate, and maximum allowable force (e.g., gentler and slower for soft/fragile objects; faster and firmer for hard objects), while using confidence thresholds to postpone adaptation when uncertainty is high.
>
> **Shared control and personalization:** mappings from tactile variables to control actions can be tuned per user and task context to reduce cognitive load and improve perceived naturalness, potentially leveraging user-in-the-loop calibration procedures.

These directions align with the broader goal of sensory-enabled prostheses: tactile sensing becomes most impactful when it directly improves manipulation stability and reduces the need for constant visual monitoring [9, 10, 6].

Optimizing feedback coding for usability and reduced cognitive load. The distributed feedback concept demonstrated here can be extended by investigating stimulus coding strategies that maximize interpretability while remaining comfortable and unobtrusive. This includes studying spatial coding (which sensors map to which stimulation sites), temporal coding (pulse trains, event-driven stimulation), and intensity coding (how amplitude or duty cycle reflects tactile magnitude), guided by wearable haptics and tactile display principles [13, 14]. In addition, confidence-aware feedback strategies become possible based on the probabilistic outputs studied in Chapter 3: the system can choose to provide stronger, weaker, or delayed feedback depending on uncertainty, thereby avoiding misleading sensations and improving user trust.

Evaluation beyond the laboratory: amputee studies and longitudinal validation. A major future milestone is to validate the proposed framework with people with limb loss, and under longer-term usage conditions. This includes assessing functional outcomes (task performance, grasp stability), cognitive load, embodiment, comfort, and learning effects over time [210, 9]. Such studies are essential to determine what tactile variables are most useful in daily life, how users adapt to feedback, and what trade-offs between bandwidth and usability are acceptable.

Toward embedded and wearable implementations. While parts of the real-time pipeline were demonstrated on a PC, and the sensing electronics were designed for practical acquisition, a natural direction is to move more computation closer to the prosthesis. Future work may target embedded implementations of key processing stages (filtering, event detection, lightweight inference, and feedback mapping) with strict constraints on power and latency. This would enable fully portable, self-contained systems suitable for everyday use, and would align with the broader trend toward edge intelligence in tactile sensing systems.

Overall, the future directions above aim to turn the thesis contributions into a more complete sensory-enabled prosthesis framework: robust distributed sensorization, generalizable online tactile inference, feedback coding optimized for human use, and tactile-informed control policies. Advancing along these axes will move tactile prosthetic systems closer to the long-term goal of restoring a useful and intuitive sense of touch in upper-limb prostheses.

Acknowledgements

This thesis represents the culmination of three years of hard work and dedication, and it is with deep gratitude that I acknowledge the many people who have supported me along the way.

Firstly, I extend my heartfelt thanks to my supervisor, Prof. Maurizio Valle, for his invaluable guidance, insightful feedback, and unwavering encouragement. His expertise and dedication have been crucial in shaping the direction and quality of my research. I also wish to acknowledge Dr. Yahya Abbass, and Dr. Nicolò Boccardo for their continued guidance and feedback.

I am particularly grateful to Prof. Strahinja Dosen from Aalborg University for his support and supervision and hosting during my period abroad. I also want to express my gratitude to the University of Genova in Italy for fostering a positive working environment through their dedicated staff.

Hussein Bassal

Publication record

1. International Journal Papers

[1] Bassal, Hussein, Yahya Abbass, Christian Gianoglio, and Maurizio Valle. "Hardness discrimination using piezoelectric-based biomimetic tactile sensor and machine learning." IEEE Sensors Letters (2024).

[2] Bassal, Hussein, Yahya Abbass, Christian Gianoglio, and Maurizio Valle. "Piezoelectric Sensor behavior analysis for Hardness Classification." To be submitted to Q1 Journal e.g.

[3] Bassal, Hussein, Nicolò Boccardo, Yahya Abbass, and Maurizio Valle. "Real-time Hardness Classification using Hannes Prosthetic hand and Machine Learning." To be submitted to Q1 Journal e.g.

[4] Bassal, Hussein, Yahya Abbass, Strahinja Dosen, Nicolò Boccardo, and Maurizio Valle. "End-to-End Electrotactile Feedback system for Anthropomorphic Prosthetic Hands." To be submitted to Q1 Journal e.g.

2. International Conference Papers

[5] Bassal, Hussein, Nicolò Boccardo, Yahya Abbass, Matteo Laffranchi, and Maurizio Valle. "Biomimetic Tactile Sensing for Hannes Anthropomorphic Prosthetic Hand." In 2024 31st IEEE International Conference on Electronics, Circuits and Systems (ICECS), pp. 1-4. IEEE, 2024.

[6] Bassal, H., Y. Abbass, and M. Valle. "Enhancing Motor Control and Feedback in Prosthetic Hands Through Piezoelectric Tactile Sensors: A Comprehensive Methodological Approach." In Summer School on Neurorehabilitation, pp. 15-18. Cham: Springer Nature Switzerland, 2022.

Participation in Research Projects

Research project 1, EU project IntelliMan

This project is focusing on the question of “How a robot can efficiently learn to manipulate in a purposeful and highly performant way”. IntelliMan will range from learning individual manipulation skills from human demonstration, to learning abstract descriptions of a manipulation task suitable for high-level planning, to discovering an object’s functionality by interacting with it, to guarantee performance and safety. IntelliMan aims at developing a novel AI-Powered Manipulation System with persistent learning capabilities, able to perceive the main characteristics and features of its surrounding by means of a heterogeneous set of sensors, able to decide how to execute a task in an autonomous way and able to detect failures in the task execution in order to request new knowledge through the interaction with humans and the environment. IntelliMan further investigates how such AI-powered manipulation systems are perceived by the users and what factors enhance human acceptability.

Activity: My developed research during the PhD was a contribution for the IntelliMan project, and this contribution could be summarized under these bullet points:

Developed a tactile sensing system for prosthetic hands.

Developed real-time pipeline for data processing and classification (Cartesian / Hannes).

Teaching and other activities

Reviewing activity for the IEEE International Symposium on Circuits and Systems ISCAS 2026

Period abroad PhD Visitor: During the PhD I had the opportunity to be a visiting PhD guest at Istituto Italiano di Tecnologia at Rehab Lab starting from November 2023 till July 2024 where I developed the real time pipeline of hardness classification using piezoelectric sensors and Hannes hand. Later on I moved to Aalborg University as PhD guest starting November 2024 till October 2025 where I developed the End-to-End Electrotactile feedback system.

Bibliography

[1] L. Derksen, "Visualising high-dimensional datasets using PCA and t-SNE in Python," 2019, medium.

[2] K. Lee, T. Ikeda, T. Miyashita, H. Ishiguro, and N. Hagita, "Separation of tactile information from multiple sources based on spatial ICA and time series clustering," in *2011 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2011, pp. 791–796.

[3] P. Xanthopoulos, P. M. Pardalos, and T. B. Trafalis, "Linear discriminant analysis," in *Robust Data Mining*. Springer, 2013, pp. 27–33.

[4] B. A. Jenkins and E. A. Lumpkin, "Developing a sense of touch," *Development*, vol. 144, no. 22, pp. 4078–4090, 2017.

[5] K. O. Johnson, T. Yoshioka, and F. Vega–Bermudez, "Tactile Functions of Mechanoreceptive Afferents Innervating the Hand," *Journal of Clinical Neurophysiology*, vol. 17, no. 6, pp. 539–558, nov 2000. [Online]. Available: http://journals.lww.com/00004691-200011000-00002

[6] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks," *Nature Reviews Neuroscience*, vol. 10, no. 5, pp. 345–359, 2009.

[7] P. D. Marasco, K. Kim, J. E. Colgate, M. A. Peshkin, and T. A. Kuiken, "Robotic touch shifts perception of embodiment to a prosthesis in targeted reinnervation amputees," *Brain*, vol. 134, no. 3, pp. 747–758, mar 2011. [Online]. Available: https://academic.oup.com/brain/article-lookup/doi/10.1093/brain/awq361

[8] S. J. Bensmaia, L. E. Miller, and S. Micera, "Restoring sensory information via bionic hands," *Nature Biomedical Engineering*, 2020.

[9] J. S. Schofield, K. R. Evans, J. P. Carey, and J. S. Hebert, “Applications of sensory feedback in motorized upper extremity prosthesis: a review,” *Expert Review of Medical Devices*, vol. 11, no. 5, pp. 499–511, sep 2014. [Online]. Available: http://www.tandfonline.com/doi/full/10.1586/17434440.2014.929496

[10] B. Stephens-Fripp, G. Alici, and R. Mutlu, “A Review of Non-Invasive Sensory Feedback Methods for Transradial Prosthetic Hands,” *IEEE Access*, vol. 6, pp. 6878–6899, 2018. [Online]. Available: http://ieeexplore.ieee.org/document/8253455/

[11] S. Lewis, M. F. Russold, H. Dietl, and E. Kaniusas, “User demands for sensory feedback in upper extremity prostheses,” *MeMeA 2012 - 2012 IEEE Symposium on Medical Measurements and Applications, Proceedings*, pp. 188–191, 2012.

[12] B. E and C. T, “Upper-limb prosthetics: critical factors in device abandonment,” *American journal of physical medicine rehabilitation*, vol. 86, no. 12, pp. 977–987, dec 2007. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/18090439/

[13] V. Chouvardas, A. Miliou, and M. Hatalis, “Tactile displays: Overview and recent advances,” *Displays*, vol. 29, no. 3, pp. 185–194, jul 2008. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0141938207000613

[14] C. Pacchierotti, S. Sinclair, M. Solazzi, A. Frisoli, V. Hayward, and D. Prattichizzo, “Wearable Haptic Systems for the Fingertip and the Hand: Taxonomy, Review, and Perspectives,” *IEEE Transactions on Haptics*, vol. 10, no. 4, pp. 580–600, oct 2017. [Online]. Available: https://ieeexplore.ieee.org/document/7922602/

[15] M. Schiefer, D. Tan, S. M. Sidek, and D. J. Tyler, “Sensory feedback by peripheral nerve stimulation improves task performance in individuals with upper limb loss using a myoelectric prosthesis,” *Journal of Neural Engineering*, vol. 13, no. 1, pp. 1–25, 2015.

[16] K. Li, Y. Fang, Y. Zhou, and H. Liu, “Non-Invasive Stimulation-Based Tactile Sensation for Upper-Extremity Prosthesis: A Review,” *IEEE Sensors Journal*, vol. 17, no. 9, pp. 2625–2635, may 2017. [Online]. Available: https://ieeexplore.ieee.org/document/7865898/

[17] R. S. Dahiya, D. Cattin, A. Adami, C. Collini, L. Barboni, M. Valle, L. Lorenzelli, R. Oboe, S. Member, G. Metta, and F. Brunetti, “Towards Tactile Sensing System on Chip for Robotic Applications,” *IEEE Sensors Journal*, vol. 11, no. 12, pp. 3216–3226, 2011.

[18] R. S. Dahiya, P. Mittendorfer, M. Valle, G. Cheng, and V. J. Lumelsky, “Directions Toward Effective Utilization of Tactile Skin: A Review,” *IEEE Sensors Journal*, vol. 13, no. 11, pp. 4121–4138, nov 2013. [Online]. Available: http://ieeexplore.ieee.org/document/6583342/

[19] S. Li, Y. Zhang, Y. Wang, K. Xia, Z. Yin, H. Wang, M. Zhang, X. Liang, H. Lu, M. Zhu, H. Wang, X. Shen, and Y. Zhang, “Physical sensors for skin-inspired electronics,” *InfoMat*, vol. 2, no. 1, pp. 184–211, 2020.

[20] N. Yogeswaran, W. Dang, W. Navaraj, D. Shakthivel, S. Khan, E. Polat, S. Gupta, H. Heidari, M. Kaboli, L. Lorenzelli, G. Cheng, and R. Dahiya, “New materials and advances in making electronic skin for interactive robots,” *Advanced Robotics*, vol. 29, no. 21, pp. 1359–1373, nov 2015. [Online]. Available: http://www.tandfonline.com/doi/full/10.1080/01691864.2015.1095653

[21] A. Nag, S. C. Mukhopadhyay, and J. Kosel, “Wearable Flexible Sensors: A Review,” *IEEE Sensors Journal*, vol. 17, no. 13, pp. 3949–3960, jul 2017. [Online]. Available: http://ieeexplore.ieee.org/document/7931559/

[22] L. Zou, C. Ge, Z. Wang, E. Cretu, and X. Li, “Novel Tactile Sensor Technology and Smart Tactile Sensing Systems: A Review,” *Sensors*, vol. 17, no. 11, p. 2653, nov 2017. [Online]. Available: http://www.mdpi.com/1424-8220/17/11/2653

[23] M. I. Tiwana, S. J. Redmond, and N. H. Lovell, “A review of tactile sensing technologies with applications in biomedical engineering,” *Sensors and Actuators A: Physical*, vol. 179, pp. 17–31, jun 2012. [Online]. Available: http://dx.doi.org/10.1016/j.sna.2012.02.051https://linkinghub.elsevier.com/retrieve/pii/S0924424712001641

[24] H. Yousef, M. Boukallel, and K. Althoefer, “Tactile sensing for dexterous in-hand manipulation in robotics—A review,” *Sensors and Actuators A: Physical*, vol. 167,

no. 2, pp. 171–187, jun 2011. [Online]. Available: http://dx.doi.org/10.1016/j.sna.2011.02.038https://linkinghub.elsevier.com/retrieve/pii/S0924424711001105

[25] P. Roberts, M. Zadan, and C. Majidi, "Soft Tactile Sensing Skins for Robotics," *Current Robotics Reports*, vol. 2, no. 3, pp. 343–354, sep 2021. [Online]. Available: https://link.springer.com/10.1007/s43154-021-00065-2

[26] L. Pinna, A. Ibrahim, and M. Valle, "Interface Electronics for Tactile Sensors Based on Piezoelectric Polymers," *IEEE Sensors Journal*, vol. 17, no. 18, pp. 5937–5947, sep 2017. [Online]. Available: http://ieeexplore.ieee.org/document/7990126/

[27] M. Saleh, Y. Abbass, A. Ibrahim, and M. Valle, "Experimental Assessment of the Interface Electronic System for PVDF-Based Piezoelectric Tactile Sensors," *Sensors (Switzerland)*, vol. 19, no. 20, pp. 1–12, 2019.

[28] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, "Robotic tactile perception of object properties: A review," *Mechatronics*, vol. 48, no. May, pp. 54–67, dec 2017. [Online]. Available: http://dx.doi.org/10.1016/j.mechatronics.2017.11.002https://linkinghub.elsevier.com/retrieve/pii/S0957415817301575

[29] W. Chen, H. Khamis, I. Birznieks, N. F. Lepora, and S. J. Redmond, "Tactile Sensors for Friction Estimation and Incipient Slip Detection—Toward Dexterous Robotic Manipulation: A Review," *IEEE Sensors Journal*, vol. 18, no. 22, pp. 9049–9064, nov 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8453829/

[30] X. Jiang, R. Chen, and H. Zhu, "Recent progress in wearable tactile sensors combined with algorithms based on machine learning and signal processing," *APL Materials*, vol. 9, no. 3, p. 030906, 2021.

[31] K. Higashi, S. Okamoto, Y. Yamada, H. Nagano, and M. Konyo, "Hardness perception through tapping: peak and impulse of the reaction force reflect the subjective hardness," in *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*. Springer, 2018, pp. 366–375.

[32] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson, "Shape-independent hardness estimation using deep learning and a gelsight tactile sensor," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 951–958.

[33] M. Laffranchi, N. Boccardo, S. Traverso, L. Lombardi, M. Canepa, A. Lince, M. Semprini, J. A. Saglia, A. Naceri, R. Sacchetti, E. Gruppioni, and L. De Michieli, "The Hannes hand prosthesis replicates the key biological properties of the human hand," *Science Robotics*, vol. 5, no. 46, p. 467, sep 2020. [Online]. Available: https://www.science.org/doi/10.1126/scirobotics.abb0467

[34] D. J. Tyler, "Neural interfaces for somatosensory feedback," *Current Opinion in Neurology*, vol. 28, no. 6, pp. 574–581, dec 2015. [Online]. Available: https://journals.lww.com/00019052-201512000-00005

[35] H. H. Ehrsson, B. Rosen, A. Stockselius, C. Ragno, P. Kohler, and G. Lundborg, "Upper limb amputees can be induced to experience a rubber hand as their own," *Brain*, vol. 131, no. 12, pp. 3443–3452, dec 2008. [Online]. Available: https://academic.oup.com/brain/article-lookup/doi/10.1093/brain/awn297

[36] V. E. Abraira and D. D. Ginty, "The sensory neurons of touch," *Neuron*, vol. 79, no. 4, pp. 618–639, 2013.

[37] S. C. Gandevia and G. Macefield, "Projection of low-threshold afferents from human intercostal muscles to the cerebral cortex," *Respiration Physiology*, vol. 77, no. 2, pp. 203–214, 1989.

[38] S. Khan, L. Lorenzelli, and R. S. Dahiya, "Technologies for printing sensors and electronics over large flexible substrates: A review," *IEEE Sensors Journal*, vol. 15, no. 6, pp. 3164–3185, 2015.

[39] R. S. Dahiya and M. Valle, *Robotic Tactile Sensing: Technologies and System*. Springer, 2013, vol. 1.

[40] Y. Abbass, M. Saleh, S. Dosen, and M. Valle, "Embedded Electrotactile Feedback System for Hand Prostheses Using Matrix Electrode and Electronic Skin," *IEEE*

Transactions on Biomedical Circuits and Systems, vol. 15, no. 5, pp. 912–925, oct 2021. [Online]. Available: https://ieeexplore.ieee.org/document/9522029/

[41] H. Fares, Y. Abbass, M. Valle, and L. Seminara, "Validation of Screen-Printed Electronic Skin Based on Piezoelectric Polymer Sensors," *Sensors*, vol. 20, no. 4, p. 1160, feb 2020. [Online]. Available: https://www.mdpi.com/1424-8220/20/4/1160

[42] S. Chen, Z. Zhou, K. Hou, X. Wu, Q. He, C. G. Tang, T. Li, X. Zhang, J. Jie, Z. Gao *et al.*, "Artificial organic afferent nerves enable closed-loop tactile feedback for intelligent robot," *Nature Communications*, vol. 15, no. 1, p. 7056, 2024.

[43] M. Alameh, Y. Abbass, A. Ibrahim, G. Moser, and M. Valle, "Touch Modality Classification Using Recurrent Neural Networks," *IEEE Sensors Journal*, vol. 21, no. 8, pp. 9983–9993, apr 2021. [Online]. Available: https://ieeexplore.ieee.org/document/9340346/

[44] A. Ibrahim, L. Pinna, L. Seminara, and M. Valle, "Achievements and Open Issues Toward Embedding Tactile Sensing and Interpretation into Electronic Skin Systems," in *Material-Integrated Intelligent Systems - Technology and Applications*. Weinheim, Germany: Wiley-VCH Verlag GmbH Co. KGaA, dec 2017, pp. 571–594. [Online]. Available: http://doi.wiley.com/10.1002/9783527679249.ch23https://onlinelibrary.wiley.com/doi/10.1002/9783527679249.ch23

[45] R. Dahiya, G. Metta, M. Valle, and G. Sandini, "Tactile Sensing—From Humans to Humanoids," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 1–20, feb 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5339133/

[46] L. Seminara, M. Capurro, P. Cirillo, G. Cannata, and M. Valle, "Electromechanical characterization of piezoelectric PVDF polymer films for tactile sensors in robotics applications," *Sensors and Actuators A: Physical*, vol. 169, no. 1, pp. 49–58, sep 2011. [Online]. Available: http://dx.doi.org/10.1016/j.sna.2011.05.004https://linkinghub.elsevier.com/retrieve/pii/S0924424711003220

[47] H. Muhammad, C. Oddo, L. Beccai, C. Recchiuto, C. Anthony, M. Adams, M. Carrozza, D. Hukins, and M. Ward, "Development of

a bioinspired MEMS based capacitive tactile sensor for a robotic finger," *Sensors and Actuators A: Physical*, vol. 165, no. 2, pp. 221–229, feb 2011. [Online]. Available: http://dx.doi.org/10.1016/j.sna.2010.10.025https://linkinghub.elsevier.com/retrieve/pii/S0924424710004838

[48] D. P. Cotton, P. H. Chappell, A. Cranny, N. M. White, and S. P. Beeby, "A novel thick-film piezoelectric slip sensor for a prosthetic hand," *IEEE Sensors Journal*, vol. 7, no. 5, pp. 752–761, 2007.

[49] I. Dakua and N. Afzulpurkar, "Piezoelectric Energy Generation and Harvesting at the Nano-Scale: Materials and Devices," *Nanomaterials and Nanotechnology*, vol. 3, no. 1, p. 21, jun 2013. [Online]. Available: http://journals.sagepub.com/doi/10.5772/56941

[50] Shinya Takamuku, G. Gomez, Koh Hosoda, and R. Pfeifer, "Haptic discrimination of material properties by a robotic hand," in *2007 IEEE 6th International Conference on Development and Learning*. IEEE, jul 2007, pp. 1–6. [Online]. Available: http://ieeexplore.ieee.org/document/4354057/

[51] C. M. Oddo, M. Controzzi, L. Beccai, C. Cipriani, and M. C. Carrozza, "Roughness Encoding for Discrimination of Surfaces in Artificial Active-Touch," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 522–533, jun 2011. [Online]. Available: http://ieeexplore.ieee.org/document/5740988/

[52] W. Liu, P. Yu, C. Gu, X. Cheng, and X. Fu, "Fingertip Piezoelectric Tactile Sensor Array for Roughness Encoding Under Varying Scanning Velocity," *IEEE Sensors Journal*, vol. 17, no. 21, pp. 6867–6879, nov 2017. [Online]. Available: http://ieeexplore.ieee.org/document/7962146/

[53] S. Omata and Y. Terunuma, "New tactile sensor like the human hand and its applications," *Sensors and Actuators A: Physical*, vol. 35, no. 1, pp. 9–15, oct 1992. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/092442479287002X

[54] B. C.-K. Tee, A. Chortos, A. Berndt, A. K. Nguyen, A. Tom, A. McGuire, Z. C. Lin, K. Tien, W.-G. Bae, H. Wang, P. Mei, H.-H. Chou, B. Cui, K. Deisseroth,

T. N. Ng, and Z. Bao, "A skin-inspired organic digital mechanoreceptor," *Science*, vol. 350, no. 6258, pp. 313–316, oct 2015. [Online]. Available: https://www.sciencemag.org/lookup/doi/10.1126/science.aaa9306

[55] A. Chortos, J. Liu, and Z. Bao, "Pursuing prosthetic electronic skin," *Nature Materials*, vol. 15, no. 9, pp. 937–950, sep 2016. [Online]. Available: http://www.nature.com/articles/nmat4671

[56] N. Jorgovanovic, S. Dosen, D. J. Djozic, G. Krajoski, and D. Farina, "Virtual Grasping: Closed-Loop Force Control Using Electrotactile Feedback," *Computational and Mathematical Methods in Medicine*, vol. 2014, pp. 1–13, 2014. [Online]. Available: https://www.hindawi.com/journals/cmmm/2014/120357/

[57] B. J. Kane, M. R. Cutkosky, and G. T. Kovacs, "CMOS-compatible traction stress sensor for use in high-resolution tactile imaging," *Sensors and Actuators A: Physical*, vol. 54, no. 1-3, pp. 511–516, jun 1996. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0924424795011919

[58] M. Kalantari, M. Ramezanifard, R. Ahmadi, J. Dargahi, and J. Kövecses, "A piezoresistive tactile sensor for tissue characterization during catheter-based cardiac surgery," *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 7, no. 4, pp. 431–440, dec 2011. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/rcs.413

[59] R. Koiva, M. Zenker, C. Schurmann, R. Haschke, and H. J. Ritter, "A highly sensitive 3D-shaped tactile sensor," in *2013 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. IEEE, jul 2013, pp. 1084–1089. [Online]. Available: http://ieeexplore.ieee.org/document/6584238/

[60] H. Muhammad, C. Recchiuto, C. Oddo, L. Beccai, C. Anthony, M. Adams, M. Carrozza, and M. Ward, "A capacitive tactile sensor array for surface texture discrimination," *Microelectronic Engineering*, vol. 88, no. 8, pp. 1811–1813, aug 2011. [Online]. Available: http://dx.doi.org/10.1016/j.mee.2011.01.045https://linkinghub.elsevier.com/retrieve/pii/S0167931711000700

[61] Y. Wang, K. Xi, G. Liang, M. Mei, and Z. Chen, "A flexible capacitive tactile sensor array for prosthetic hand real-time contact force measurement," *2014 IEEE International Conference on Information and Automation, ICIA 2014*, vol. 8, no. July, pp. 937–942, 2014.

[62] H.-K. Kim, S. Lee, and K.-S. Yun, "Capacitive tactile sensor array for touch screen application," *Sensors and Actuators A: Physical*, vol. 165, no. 1, pp. 2–7, jan 2011. [Online]. Available: http://dx.doi.org/10.1016/j.sna.2009.12.031https://linkinghub.elsevier.com/retrieve/pii/S0924424709005573

[63] H. Zhao, K. O'Brien, S. Li, and R. F. Shepherd, "Optoelectronically innervated soft prosthetic hand via stretchable optical waveguides," *Science Robotics*, vol. 1, no. 1, p. eaai7529, dec 2016. [Online]. Available: https://robotics.sciencemag.org/lookup/doi/10.1126/scirobotics.aai7529

[64] A. Song, Y. Han, H. Hu, and J. Li, "A novel texture sensor for fabric texture measurement and classification," *IEEE Transactions on Instrumentation and Measurement*, vol. 63, no. 7, pp. 1739–1747, 2013.

[65] S. Chun, A. Hong, Y. Choi, C. Ha, and W. Park, "A tactile sensor using a conductive graphene-sponge composite," *Nanoscale*, vol. 8, no. 17, pp. 9185–9192, 2016.

[66] Y. Cao, T. Li, Y. Gu, H. Luo, S. Wang, and T. Zhang, "Fingerprint-inspired flexible tactile sensor for accurately discerning surface texture," *Small*, vol. 14, no. 16, p. 1703902, 2018.

[67] L. Qin, Z. Yi, and Y. Zhang, "Enhanced surface roughness discrimination with optimized features from bio-inspired tactile sensor," *Sensors and Actuators A: Physical*, vol. 264, pp. 133–140, 2017.

[68] O. Kursun and A. Patooghy, "An embedded system for collection and real-time classification of a tactile dataset," *IEEE Access*, vol. 8, pp. 97 462–97 473, 2020.

[69] L. Hertel, H. Phan, and A. Mertins, "Comparing time and frequency domain for audio event recognition using deep learning," in *2016 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2016, pp. 3407–3411.

[70] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.

[71] A. Shrestha and A. Mahmood, "Review of deep learning algorithms and architectures," *IEEE Access*, vol. 7, pp. 53 040–53 065, 2019.

[72] S. A. A. Karim, M. H. Kamarudin, B. A. Karim, M. K. Hasan, and J. Sulaiman, "Wavelet transform and fast fourier transform for signal compression: A comparative study," in *2011 International Conference on Electronic Devices, Systems and Applications (ICEDSA)*. IEEE, 2011, pp. 280–285.

[73] S. Mallat, *A Wavelet Tour of Signal Processing*. Elsevier, 1999.

[74] N. M. Hanif, P. H. Chappell, A. Cranny, and N. M. White, "Surface texture detection with artificial fingers," in *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2015, pp. 8018–8021.

[75] Z. Zhang, J. Zhou, Z. Yan, K. Wang, J. Mao, and Z. Jiang, "Hardness recognition of fruits and vegetables based on tactile array information of manipulator," *Computers and Electronics in Agriculture*, vol. 181, p. 105959, 2021.

[76] F. De Boissieu, C. Godin, B. Guilhamat, D. David, C. Serviere, and D. Baudois, "Tactile texture recognition with a 3-axial force MEMS integrated artificial finger," in *Robotics: Science and Systems*, Seattle, WA, 2009, pp. 49–56.

[77] A. Ibrahim, L. Noli, H. Chible, and M. Valle, "Embedded electronic systems for tactile data processing," in *International Conference on Applications in Electronics Pervading Industry, Environment and Society*. Springer, 2016, pp. 17–23.

[78] S. Wold, K. Esbensen, and P. Geladi, "Principal component analysis," *Chemometrics and Intelligent Laboratory Systems*, vol. 2, no. 1-3, pp. 37–52, 1987.

[79] E. Kerr, T. McGinnity, and S. Coleman, "Material recognition using tactile sensing," *Expert Systems with Applications*, vol. 94, pp. 94–111, 2018.

[80] E. Kerr, T. M. McGinnity, and S. Coleman, “Material classification based on thermal properties—a robot and human evaluation,” in *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2013, pp. 1048–1053.

[81] H. Hu, Y. Han, A. Song, S. Chen, C. Wang, and Z. Wang, “A finger-shaped tactile sensor for fabric surfaces evaluation by 2-dimensional active sliding touch,” *Sensors*, vol. 14, no. 3, pp. 4899–4913, 2014.

[82] B. Schölkopf, A. Smola, and K.-R. Müller, “Nonlinear component analysis as a kernel eigenvalue problem,” *Neural Computation*, vol. 10, no. 5, pp. 1299–1319, 1998.

[83] Y.-H. Liu, Y.-T. Hsiao, W.-T. Cheng, Y.-C. Liu, and J.-Y. Su, “Low-resolution tactile image recognition for automated robotic assembly using kernel PCA-based feature fusion and multiple kernel learning-based support vector machine,” *Mathematical Problems in Engineering*, vol. 2014, 2014.

[84] A. Weingessel and K. Hornik, “Local PCA algorithms,” *IEEE Transactions on Neural Networks*, vol. 11, no. 6, pp. 1242–1250, 2000.

[85] M. Schopfer, H. Ritter, and G. Heidemann, “Acquisition and application of a tactile database,” in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. IEEE, 2007, pp. 1517–1522.

[86] P. Comon, “Independent component analysis, a new concept?” *Signal Processing*, vol. 36, no. 3, pp. 287–314, 1994.

[87] J. V. Stone, “Independent component analysis: an introduction,” *Trends in Cognitive Sciences*, vol. 6, no. 2, pp. 59–64, 2002.

[88] M. Pal, A. Khasnobish, A. Konar, D. Tibarewala, and R. Janarthanan, “Classification of deformable and non-deformable surfaces by tactile image analysis,” in *Proceedings of The 2014 International Conference on Control, Instrumentation, Energy and Communication (CIEC)*. IEEE, 2014, pp. 626–630.

[89] H. Nguyen, L. Osborn, M. Iskarous, C. Shallal, C. Hunt, J. Betthauser, and N. Thakor, “Dynamic texture decoding using a neuromorphic multilayer tactile

sensor," in *2018 IEEE Biomedical Circuits and Systems Conference (BioCAS)*. IEEE, 2018, pp. 1–4.

[90] J. Hoelscher, J. Peters, and T. Hermans, "Evaluation of tactile feature extraction for interactive object recognition," in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 310–317.

[91] D. H. Wolpert and W. G. Macready, "No free lunch theorems for optimization," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 67–82, Apr. 1997.

[92] B. Mahesh, "Machine learning algorithms - a review," *International Journal of Science and Research (IJSR)*, vol. 9, no. 1, pp. 381–386, 2020.

[93] A. Masteller, S. Sankar, H. B. Kim, K. Ding, X. Liu, and A. H. All, "Recent Developments in Prosthesis Sensors, Texture Recognition, and Sensory Stimulation for Upper Limb Prostheses," *Annals of Biomedical Engineering*, vol. 49, no. 1, pp. 57–74, jan 2021. [Online]. Available: https://doi.org/10.1007/s10439-020-02678-8http://link.springer.com/10.1007/s10439-020-02678-8

[94] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Design of a flexible tactile sensor for classification of rigid and deformable objects," *Robotics and Autonomous Systems*, vol. 62, no. 1, pp. 3–15, 2014.

[95] W. Yuan, C. Zhu, A. Owens, M. A. Srinivasan, and E. H. Adelson, "Shape-independent hardness estimation using deep learning and a GelSight tactile sensor," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 951–958.

[96] I. Bandyopadhyaya, D. Babu, A. Kumar, and J. Roychowdhury, "Tactile sensing based softness classification using machine learning," in *2014 IEEE International Advance Computing Conference (IACC)*. IEEE, feb 2014, pp. 1231–1236. [Online]. Available: http://ieeexplore.ieee.org/document/6779503/

[97] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, “Haptic classification and recognition of objects using a tactile sensing forearm,” in *2012 IEEE/RSJ international Conference on intelligent Robots and systems*. IEEE, 2012, pp. 4090–4097.

[98] İ. Bogrekcı, P. Demırcıoglu, H. S. Sucuoglu, and O. Turhanlar, “The effect of the infill type and density on hardness of 3d printed parts,” *International journal of 3d printing technologies and digital industry*, vol. 3, no. 3, pp. 212–219, 2019.

[99] D. S. Chathuranga, S. Hirai *et al.*, “Investigation of a biomimetic fingertip’s ability to discriminate fabrics based on surface textures,” in *2013 IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. IEEE, 2013, pp. 1667–1674.

[100] T. Yongting, Z. Jun, M. Yimeng, Z. Na, and Y. Xiaorong, “Design and experiment of tactile sensors for testing surface roughness of fruits and vegetables,” *Nongye Jixie Xuebao/Transactions of the Chinese Society of Agricultural Machinery*, vol. 46, no. 11, 2015.

[101] T. Araki, M. Makikawa, S. Hirai *et al.*, “Experimental investigation of surface identification ability of a low-profile fabric tactile sensor,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 4497–4504.

[102] M. Kaboli, P. Mittendorfer, V. Hügel, and G. Cheng, “Humanoids learn object properties from robust tactile feature descriptors via multi-modal artificial skin,” in *2014 IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2014, pp. 187–192.

[103] A. Drimus, M. B. Petersen, and A. Bilberg, “Object texture recognition by dynamic tactile sensing using active exploration,” in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 2012, pp. 277–283.

[104] N. Jamali and C. Sammut, “Material classification by tactile sensing using surface textures,” in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2336–2341.

[105] M. Kaboli, R. Walker, and G. Cheng, “Re-using prior tactile experience by robotic hands to discriminate in-hand objects via texture properties,” in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 2242–2247.

[106] S. Luo, W. Mou, M. Li, K. Althoefer, and H. Liu, “Rotation and translation invariant object recognition with a tactile sensor,” in *SENSORS, 2014 IEEE*. IEEE, 2014, pp. 1030–1033.

[107] Z. Su, K. Hausman, Y. Chebotar, A. Molchanov, G. E. Loeb, G. S. Sukhatme, and S. Schaal, “Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor,” in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2015, pp. 297–303.

[108] J. Schill, J. Laaksonen, M. Przybylski, V. Kyrki, T. Asfour, and R. Dillmann, “Learning continuous grasp stability for a humanoid robot hand based on tactile sensing,” in *2012 4th IEEE RAS & EMBS international conference on biomedical robotics and biomechatronics (BioRob)*. IEEE, 2012, pp. 1901–1906.

[109] T. Li, X. Sun, X. Shu, C. Wang, Y. Wang, G. Chen, and N. Xue, “Robot grasping system and grasp stability prediction based on flexible tactile sensor array,” *Machines*, vol. 9, no. 6, p. 119, 2021.

[110] S. Funabashi, T. Isobe, S. Ogasa, T. Ogata, A. Schmitz, T. P. Tomo, and S. Sugano, “Stable in-grasp manipulation with a low-cost robot hand by using 3-axis tactile sensors with a cnn,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9166–9173.

[111] S. Funabashi, A. Schmitz, T. Sato, S. Somlor, and S. Sugano, “Versatile in-hand manipulation of objects with different sizes and shapes using neural networks,” in *2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2018, pp. 1–9.

[112] P. Gastaldo, L. Pinna, L. Seminara, M. Valle, and R. Zunino, “Computational intelligence techniques for tactile sensing systems,” *Sensors*, vol. 14, no. 6, pp. 10 952–10 976, 2014.

[113] W. M. B. Tiest, "Tactual perception of material properties," *Vision research*, vol. 50, no. 24, pp. 2775–2782, 2010.

[114] I. Bandyopadhyaya, D. Babu, A. Kumar, and J. Roychowdhury, "Tactile sensing based softness classification using machine learning," in *2014 IEEE International Advance Computing Conference (IACC)*. IEEE, 2014, pp. 1231–1236.

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

[116] H. Liu, J. Greco, X. Song, J. Bimbo, L. Seneviratne, and K. Althoefer, "Tactile image based contact shape recognition using neural network," in *2012 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*. IEEE, 2012, pp. 138–143.

[117] D. Xu, G. E. Loeb, and J. A. Fishel, "Tactile identification of objects using bayesian exploration," in *2013 IEEE international conference on robotics and automation*. IEEE, 2013, pp. 3056–3061.

[118] A. J. Silva, O. A. D. Ramirez, V. P. Vega, and J. P. O. Oliver, "Phantom omni haptic device: Kinematic and manipulability," in *2009 Electronics, Robotics and Automotive Mechanics Conference (CERMA)*. IEEE, 2009, pp. 193–198.

[119] T. Asfour, K. Regenstein, P. Azad, J. Schroder, A. Bierbaum, N. Vahrenkamp, and R. Dillmann, "Armar-iii: An integrated humanoid platform for sensory-motor control," in *2006 6th IEEE-RAS international conference on humanoid robots*. IEEE, 2006, pp. 169–175.

[120] C. Antfolk, M. D'Alonzo, B. Rosén, G. Lundborg, F. Sebelius, and C. Cipriani, "Sensory feedback in upper limb prosthetics," *Expert Review of Medical Devices*, vol. 10, no. 1, pp. 45–54, jan 2013. [Online]. Available: https://www.tandfonline.com/action/journalInformation?journalCode=ierd20http://www.tandfonline.com/doi/full/10.1586/erd.12.68

[121] P. Svensson, U. Wijk, A. Björkman, and C. Antfolk, "A review of invasive and non-invasive sensory feedback in upper limb prostheses," *Expert*

Review of Medical Devices, vol. 14, no. 6, pp. 439–447, jun 2017. [Online]. Available: http://dx.doi.org/10.1080/17434440.2017.1332989https://www.tandfonline.com/doi/full/10.1080/17434440.2017.1332989

[122] J. W. Sensinger and S. Dosen, "A review of sensory feedback in upper-limb prostheses from the perspective of human motor control," *Frontiers in Neuroscience*, vol. 14, p. 345, 2020.

[123] H. Culbertson, S. B. Schorr, and A. M. Okamura, "Haptics: The Present and Future of Artificial Touch Sensation," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 385–409, may 2018. [Online]. Available: https://www.annualreviews.org/doi/abs/10.1146/annurev-control-060117-105043https://www.annualreviews.org/doi/10.1146/annurev-control-060117-105043

[124] D. Prattichizzo, C. Pacchierotti, and G. Rosati, "Cutaneous force feedback as a sensory subtraction technique in haptics," *IEEE Transactions on Haptics*, vol. 5, no. 4, pp. 289–300, 2012.

[125] R. L. Klatzky and S. J. Lederman, "Touch," in *Handbook of Psychology*. Hoboken, NJ, USA: John Wiley Sons, Inc., apr 2003, vol. 14, no. 8, pp. 147–176. [Online]. Available: https://onlinelibrary.wiley.com/doi/full/10.1002/0471264385.wei0406https://onlinelibrary.wiley.com/doi/abs/10.1002/0471264385.wei0406https://onlinelibrary.wiley.com/doi/10.1002/0471264385.wei0406

[126] J. W. Kaczmarek, K.A.(Dept. of Electr. Eng., Wisconsin Univ., Madison, WI, USA) and U. (Dept. of Electr. Eng., Wisconsin Univ., Madison, WI, "Electrotactile and vibrotactile displays for sensory substitution systems," *IEEE Transactions on Biomedical Engineering (Volume: 38 , Issue:)*, vol. 38, no. 1, pp. 1–16, 1991.

[127] R. W. Cholewiak and A. A. Collins, "Vibrotactile localization on the arm: Effects of place, space, and age," *Perception Psychophysics*, vol. 65, no. 7, pp. 1058–1077, oct 2003. [Online]. Available: http://link.springer.com/10.3758/BF03194834

[128] M. Fleury, G. Lioi, C. Barillot, and A. Lécuyer, "A Survey on the Use of Haptic Feedback for Brain-Computer Interfaces and Neurofeedback," *Frontiers in Neuroscience*, vol. 14, no. June, pp. 1–16, jun 2020. [Online]. Available: https://www.frontiersin.org/article/10.3389/fnins.2020.00528/full

[129] "Michelangelo prosthetic hand | Ottobock US." [Online]. Available: https://www.ottobockus.com/prosthetics/upper-limb-prosthetics/solution-overview/michelangelo-prosthetic-hand/

[130] J. Pons, E. Rocon, R. Ceres, D. Reynaerts, B. Saro, S. Levin, and W. Van Moorleghem, "The MANUS-HAND Dextrous Robotics Upper Limb Prosthesis: Mechanical and Manipulation Aspects," *Autonomous Robots*, vol. 16, no. 2, pp. 143–163, mar 2004. [Online]. Available: http://link.springer.com/10.1023/B:AURO.0000016862.38337.f1

[131] A. Kargov, O. Ivlev, C. Pylatiuk, T. Asfour, S. Schulz, A. Graeser, R. Dillmann, and G. Bretthauer, "Applications of a Fluidic Artificial Hand in the Field of Rehabilitation," in *Rehabilitation Robotics*. I-Tech Education and Publishing, aug 2007, no. August. [Online]. Available: http://www.intechopen.com/books/rehabilitation_robotics/applications_of_a_fluidic_artificial_hand_in_the_field_of_rehabilitation

[132] C. Antfolk, C. Balkenius, B. Rosén, G. Lundborg, and F. Sebelius, "SmartHand tactile display: A new concept for providing sensory feedback in hand prostheses," *Journal of Plastic Surgery and Hand Surgery*, vol. 44, no. 1, pp. 50–53, 2010.

[133] L. A. Jones and N. B. Sarter, "Tactile Displays: Guidance for Their Design and Application," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 50, no. 1, pp. 90–111, feb 2008. [Online]. Available: http://journals.sagepub.com/doi/10.1518/001872008X250638

[134] K. A. Kaczmarek, "Electrotactile adaptation on the abdomen: Preliminary results," *IEEE Transactions on Rehabilitation Engineering*, vol. 8, no. 4, pp. 499–505, 2000.

[135] I. R. Summers, P. R. Dixon, P. G. Cooper, D. A. Gratton, B. H. Brown, and J. C. Stevens, "Vibrotactile and electrotactile perception of time-varying pulse

trains," *The Journal of the Acoustical Society of America*, vol. 95, no. 3, pp. 1548–1558, mar 1994. [Online]. Available: https://asa.scitation.org/doi/abs/10.1121/1.408542http://asa.scitation.org/doi/10.1121/1.408542

[136] K. Kaczmarek and S. Haase, "Pattern identification as a function of stimulation on a fingertip-scanned electrotactile display," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 11, no. 3, pp. 269–275, sep 2003. [Online]. Available: https://ieeexplore.ieee.org/document/1231236/

[137] K. Wolf and T. Bäder, "Illusion of Surface Changes Induced by Tactile and Visual Touch Feedback," in *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, vol. 18. New York, NY, USA: ACM, apr 2015, pp. 1355–1360. [Online]. Available: https://dl.acm.org/doi/10.1145/2702613.2732703

[138] A. W. Shehata, M. Rehani, Z. E. Jassat, and J. S. Hebert, "Mechanotactile sensory feedback improves embodiment of a prosthetic hand during active use," *Frontiers in Neuroscience*, vol. 14, p. 263, 2020.

[139] K. Bark, J. W. Wheeler, S. Premakumar, and M. R. Cutkosky, "Comparison of skin stretch and vibrotactile stimulation for feedback of proprioceptive information," in *2008 Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2008, pp. 71–78.

[140] J. Wheeler, K. Bark, J. Savall, and M. Cutkosky, "Investigation of rotational skin stretch for proprioceptive feedback with application to myoelectric systems," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 1, pp. 58–66, 2010.

[141] L. P. Paredes, S. Dosen, F. Rattay, B. Graimann, and D. Farina, "The impact of the stimulation frequency on closed-loop control with electrotactile feedback," *Journal of NeuroEngineering and Rehabilitation*, vol. 12, no. 35, 2015.

[142] A. R. Krueger, P. Giannoni, V. Shah, M. Casadio, and R. A. Scheidt, "Supplemental vibrotactile feedback control of stabilization and reaching actions of the arm using limb state and position error encodings," *Journal of NeuroEngineering and Rehabilitation*, vol. 14, no. 36, 2017.

[143] V. R. Borkowska, A. McConnell, S. Vijayakumar, A. A. Stokes, and A. D. Roche, “A haptic sleeve as a method of mechanotactile feedback restoration for myoelectric hand prosthesis users,” *Frontiers in Rehabilitation Sciences*, vol. 3, p. 806479, 2022.

[144] M. Laghi, M. G. Catalano, G. Grioli, and A. Bicchi, “A wearable wrist haptic display for motion tracking and force feedback in the operational space,” *Wearable Technologies*, vol. 2, p. e5, 2021.

[145] H. J. B. Witteveen, E. A. Droog, J. S. Rietman, and P. H. Veltink, “Vibro- and Electrotactile User Feedback on Hand Opening for Myoelectric Forearm Prostheses,” *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 8, pp. 2219–2226, aug 2012. [Online]. Available: http://ieeexplore.ieee.org/document/6203566/

[146] C. Pylatiuk, A. Kargov, and S. Schulz, “Design and evaluation of a low-cost force feedback system for myoelectric prosthetic hands,” *Journal of Prosthetics and Orthotics*, vol. 18, no. 2, pp. 57–61, 2006.

[147] E. Battaglia, J. P. Clark, M. Bianchi, M. G. Catalano, A. Bicchi, and M. K. O’Malley, “Skin Stretch Haptic Feedback to Convey Closure Information in Anthropomorphic, Under-Actuated Upper Limb Soft Prostheses,” *IEEE Transactions on Haptics*, vol. 12, no. 4, pp. 508–520, oct 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8709758/

[148] K. R. Schoepp, M. R. Dawson, J. S. Schofield, J. P. Carey, and J. S. Hebert, “Design and Integration of an Inexpensive Wearable Mechanotactile Feedback System for Myoelectric Prostheses,” *IEEE Journal of Translational Engineering in Health and Medicine*, vol. 6, no. April, pp. 1–11, 2018.

[149] Y. Abbass, S. Dosen, L. Seminara, and M. Valle, “Full-hand electrotactile feedback using electronic skin and matrix electrodes for high-bandwidth human–machine interfacing,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 380, no. 2228, p. 20210017, 2022.

[150] J. Dargahi and S. Najarian, “Human tactile perception as a standard for artificial tactile sensing - a review,” *International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 01, no. 01, p. 23, 2004. [Online]. Available: http://www.roboticpublications.com/index.php?page=abstract0101&id=3

[151] K. Suwanratchatamanee, M. Matsumoto, and S. Hashimoto, “Robotic Tactile Sensor System and Applications,” *IEEE Transactions on Industrial Electronics*, vol. 57, no. 3, pp. 1074–1087, mar 2010. [Online]. Available: http://ieeexplore.ieee.org/document/5229258/

[152] J. Qu, B. Mao, Z. Li, Y. Xu, K. Zhou, X. Cao, Q. Fan, M. Xu, B. Liang, H. Liu, X. Wang, and X. Wang, “Recent Progress in Advanced Tactile Sensing Technologies for Soft Grippers,” *Advanced Functional Materials*, vol. 33, no. 41, pp. 1–32, oct 2023. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/adfm.202306249

[153] S. Sundaram, “How to improve robotic touch,” *Science*, vol. 370, no. 6518, pp. 768–769, nov 2020. [Online]. Available: https://www.science.org/doi/10.1126/science.abd3643

[154] H. Liu, D. Guo, F. Sun, W. Yang, S. Furber, and T. Sun, “Embodied tactile perception and learning,” *Brain Science Advances*, vol. 6, no. 2, pp. 132–158, jun 2020. [Online]. Available: http://journals.sagepub.com/doi/10.26599/BSA.2020.9050012

[155] B. Shih, D. Shah, J. Li, T. G. Thuruthel, Y.-L. Park, F. Iida, Z. Bao, R. Kramer-Bottiglio, and M. T. Tolley, “Electronic skins and machine learning for intelligent soft robots,” *Science Robotics*, vol. 5, no. 41, p. 9239, apr 2020. [Online]. Available: http://robotics.sciencemag.org/https://www.science.org/doi/10.1126/scirobotics.aaz9239

[156] C. Wang, C. Liu, F. Shang, S. Niu, L. Ke, N. Zhang, B. Ma, R. Li, X. Sun, and S. Zhang, “Tactile sensing technology in bionic skin: A review,” *Biosensors and Bioelectronics*, vol. 220, no. December 2021, p. 114882, jan 2023. [Online]. Available: https://doi.org/10.1016/j.bios.2022.114882https://linkinghub.elsevier.com/retrieve/pii/S0956566322009228

[157] Y. Tang, H. Liu, J. Pan, Z. Zhang, Y. Xu, N. Yao, L. Zhang, and L. Tong, "Optical Micro/Nanofiber-Enabled Compact Tactile Sensor for Hardness Discrimination," *ACS Applied Materials Interfaces*, vol. 13, no. 3, pp. 4560–4566, jan 2021. [Online]. Available: https://pubs.acs.org/doi/10.1021/acsami.0c20392

[158] Y. Tian, Y. Li, Y. Bai, L. Yang, J. Bai, F. Sun, Y. Wang, Y. Zhao, T. Li, and T. Zhang, "Development of flexible tactile sensing arrays for hardness recognition," *Sensors and Actuators A: Physical*, vol. 359, no. June, p. 114478, 2023. [Online]. Available: https://doi.org/10.1016/j.sna.2023.114478

[159] Y. Xin, M. Cui, C. Liu, T. Hou, L. Liu, C. Qian, and Y. Yan, "A bionic piezoelectric tactile sensor for features recognition of object surface based on machine learning," *Review of Scientific Instruments*, vol. 92, no. 9, 2021. [Online]. Available: https://doi.org/10.1063/5.0057236

[160] S. Gao, Q. Wang, and L. Yu, "Object Recognition Based on Hardness and Texture via Modified Force-Sensitive Fingertips of a Humanoid Hand," *IEEE Sensors Letters*, vol. 7, no. 2, pp. 1–4, 2023.

[161] Z. Zhang, J. Zhou, Z. Yan, K. Wang, J. Mao, and Z. Jiang, "Hardness recognition of fruits and vegetables based on tactile array information of manipulator," *Computers and Electronics in Agriculture*, vol. 181, no. December 2020, p. 105959, 2021. [Online]. Available: https://doi.org/10.1016/j.compag.2020.105959

[162] L. Seminara, P. Gastaldo, S. J. Watt, K. F. Valyear, F. Zuher, and F. Mastrogiovanni, "Active Haptic Perception in Robots: A Review," *Frontiers in Neurorobotics*, vol. 13, no. July, pp. 1–20, jul 2019. [Online]. Available: https://www.frontiersin.org/article/10.3389/fnbot.2019.00053/full

[163] M. O. Ernst, "Learning to integrate arbitrary signals from vision and touch," *Journal of Vision*, vol. 7, no. 5, p. 7, jun 2007. [Online]. Available: http://jov.arvojournals.org/article.aspx?doi=10.1167/7.5.7

[164] Y. Pang, X. Xu, S. Chen, Y. Fang, X. Shi, Y. Deng, Z. L. Wang, and C. Cao, "Skin-inspired textile-based tactile sensors enable multifunctional sensing of

wearables and soft robots," *Nano Energy*, vol. 96, no. March, p. 107137, 2022. [Online]. Available: https://doi.org/10.1016/j.nanoen.2022.107137

[165] "Startseite | JOANNEUM RESEARCH." [Online]. Available: https://www.joanneum.at/

[166] Y. Amin, C. Gianoglio, and M. Valle, "Embedded real-time objects' hardness classification for robotic grippers," *Future Generation Computer Systems*, vol. 148, pp. 211–224, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167739X23002182

[167] ——, "Embedded real-time objects' hardness classification for robotic grippers," *Future Generation Computer Systems*, vol. 148, pp. 211–224, nov 2023. [Online]. Available: https://doi.org/10.1016/j.future.2023.06.002https://linkinghub.elsevier.com/retrieve/pii/S0167739X23002182

[168] H. Abdi and L. J. Williams, "Principal component analysis," *WIREs Computational Statistics*, vol. 2, no. 4, pp. 433–459, jul 2010. [Online]. Available: https://wires.onlinelibrary.wiley.com/doi/10.1002/wics.101

[169] E. Kerr, T. McGinnity, and S. Coleman, "Material recognition using tactile sensing," *Expert Systems with Applications*, vol. 94, pp. 94–111, mar 2018. [Online]. Available: https://doi.org/10.1016/j.eswa.2017.10.045https://linkinghub.elsevier.com/retrieve/pii/S0957417417307273

[170] K. Murakami, K. Matsuo, T. Hasegawa, and R. Kurazume, "A decision method for placement of tactile elements on a sensor glove for the recognition of grasp types," *IEEE/ASME Transactions on Mechatronics*, vol. 15, no. 1, pp. 157–162, 2010.

[171] M. Saleh, Y. Abbass, and M. Valle, "Embedded Implementation of Signal Pre-processing for Tactile Sensing System," in *Advances in System-Integrated Intelligence*. Springer, Cham, 2023, pp. 727–736. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-031-16281-7_69https://link.springer.com/10.1007/978-3-031-16281-7_69

[172] S. M. Holland, “Principal components analysis (pca),” *Department of Geology, University of Georgia, Athens, GA*, vol. 30602, p. 2501, 2008.

[173] H. Bassal, Y. Abbass, C. Gianoglio, and M. Valle, “Hardness discrimination using piezoelectric-based biomimetic tactile sensor and machine learning,” *IEEE Sensors Letters*, 2024.

[174] I. Guyon, J. Weston, S. Barnhill, and V. Vapnik, “Gene selection for cancer classification using support vector machines,” *Machine learning*, vol. 46, no. 1, pp. 389–422, 2002.

[175] M. L. Ali, “Edge ai: Deploying machine learning models on edge devices edge ai: Deploying machine learning models on edge devices,” 06 2024.

[176] R. Hadidi, J. Cao, Y. Xie, B. Asgari, T. Krishna, and H. Kim, “Characterizing the deployment of deep neural networks on commercial edge devices,” in *2019 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 2019, pp. 35–48.

[177] A. Sawada, T. Miyagawa, A. Ebihara, S. Yachida, and T. Hosoi, “Convolutional neural networks for time-dependent classification of variable-length time series,” 07 2022, pp. 1–8.

[178] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. New York: Springer, 2009.

[179] M. Kuhn and K. Johnson, *Applied Predictive Modeling*. New York: Springer, 2013.

[180] J. Demšar, “Statistical comparisons of classifiers over multiple data sets,” *Journal of Machine Learning Research*, vol. 7, pp. 1–30, 2006.

[181] T. G. Dietterich, “Approximate statistical tests for comparing supervised classification learning algorithms,” *Neural Computation*, vol. 10, no. 7, pp. 1895–1923, 1998.

[182] S. García and F. Herrera, “An extension on “statistical comparisons of classifiers over multiple data sets” for all pairwise comparisons,” *Journal of Machine Learning Research*, vol. 9, pp. 2677–2694, 2008.

[183] B. Zhou, C. Wang, Z. Huan, Z. Li, Y. Chen, G. Gao, H. Li, C. Dong, and J. Liang, “A novel segmentation scheme with multi-probability threshold for human activity recognition using wearable sensors,” *Sensors*, vol. 22, no. 19, p. 7446, 2022.

[184] P. McCullagh and J. A. Nelder, *Generalized Linear Models*, 2nd ed. London: Chapman and Hall, 1989.

[185] S. Fang, Z. Yi, T. Mi, Z. Zhou, C. Ye, W. Shang, T. Xu, and X. Wu, “TactONet: Tactile ordinal network based on unimodal probability for object hardness classification,” *IEEE Transactions on Automation Science and Engineering*, pp. 2784–2794, 2023.

[186] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On calibration of modern neural networks,” in *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017, pp. 1321–1330.

[187] R. S. Dahiya, G. Metta, M. Valle, and G. Sandini, “Tactile sensing—from humans to humanoids,” *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 1–20, 2010.

[188] Z. Kappassov, J.-A. Corrales, and V. Perdereau, “Tactile sensing in dexterous robot hands — Review,” *Robotics and Autonomous Systems*, vol. 74, pp. 195–220, dec 2015. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0921889015001621

[189] L. Jabban, S. Dupan, D. Zhang, B. Ainsworth, K. Nazarpour, and B. Metcalfe, “Sensory feedback for upper-limb prostheses: Opportunities and barriers,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 30, pp. 738–747, 2022.

[190] A. D. Roche, Z. K. Bailey, M. Gonzalez, P. P. Vu, C. A. Chestek, D. H. Gates, S. W. P. Kemp, P. S. Cederna, M. Ortiz-Catalan, and O. C. Aszmann, “Upper limb prostheses: bridging the sensory gap,” *Journal of Hand Surgery (European Volume)*, vol. 48, no. 3, pp. 182–190, 2023.

[191] L. E. Osborn, A. Dragomir, J. L. Betthauser, C. L. Hunt, H. H. Nguyen, R. R. Kaliki, and N. V. Thakor, “Prosthesis with neuromorphic multilayered e-dermis perceives touch and pain,” *Science Robotics*, vol. 3, no. 19, p. eaat3818, jun 2018. [Online]. Available: https://robotics.sciencemag.org/lookup/doi/10.1126/scirobotics.aat3818

[192] M. Kaboli and G. Cheng, “Robust tactile descriptors for discriminating objects from textural properties via artificial robotic skin,” *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 985–1003, 2018.

[193] H. Bassal, Y. Abbass, C. Gianoglio, and M. Valle, “Hardness Discrimination using Piezoelectric-based Biomimetic Tactile Sensor and Machine Learning,” *IEEE Sensors Letters*, vol. PP, no. 3, pp. 1–4, 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10605059/

[194] N. Boccardo, M. Canepa, S. Stedman, L. Lombardi, A. Marinelli, D. Di Domenico, R. Galviati, E. Gruppioni, L. De Michieli, and M. Laffranchi, “Development of a 2-dofs actuated wrist for enhancing the dexterity of myoelectric hands,” *IEEE Transactions on Medical Robotics and Bionics*, vol. 6, no. 1, pp. 257–270, 2023.

[195] G. C. Cawley and N. L. C. Talbot, “On over-fitting in model selection and subsequent selection bias in performance evaluation,” *Journal of Machine Learning Research*, vol. 11, pp. 2079–2107, 2010.

[196] S. Raspopovic, M. Capogrosso, F. M. Petrini, M. Bonizzato, J. Rigosa, G. D. Pino, J. Carpaneto, M. Controzzi, T. Boretius, E. Fernandez, G. Granata, C. M. Oddo, L. Citi, A. L. Ciancio, C. Cipriani, M. C. Carrozza, W. Jensen, E. Guglielmelli, T. Stieglitz, P. M. Rossini, and S. Micera, “Bioengineering: Restoring natural sensory feedback in real-time bidirectional hand prostheses,” *Science Translational Medicine*, vol. 6, no. 222, pp. 1–12, 2014.

[197] S. Raspopovic, G. Valle, and F. M. Petrini, “Sensory feedback for limb prostheses in amputees,” *Nature Materials*, vol. 20, no. 7, pp. 925–939, 2021.

[198] M. Abd, A. Downey, P. P. Vu *et al.*, “Multichannel haptic feedback unlocks prosthetic hand dexterity,” *Scientific Reports*, vol. 12, p. 2244, 2022.

[199] M. Henrich, S. Povenz, C. Weninger, C. Neuper, and C. Guger, “Encoding contact size information with static and dynamic electrotactile stimulation delivered to the fingertip,” *Experimental Brain Research*, vol. 242, pp. 2049–2065, 2024.

[200] Y. Abbass, S. Dosen, L. Seminara, and M. Valle, “Full-hand electrotactile feedback using electronic skin and matrix electrodes for high-bandwidth human–machine interfacing,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 380, no. 2228, p. 20210017, 2022.

[201] M. Zirkl, A. Sawatdee, U. Helbig, M. Krause, G. Scheipl, E. Kraker, P. A. Ersman, D. Nilsson, D. Platt, P. Bodö, S. Bauer, G. Domann, and B. Stadlober, “An all-printed ferroelectric active matrix sensor network based on only five functional materials forming a touchless control interface,” *Advanced Materials*, vol. 23, no. 18, pp. 2069–2074, 2011.

[202] “32-Channel, Current-Input Analog-to-Digital Converter.” [Online]. Available: https://www.ti.com/lit/ds/symlink/ddc232.pdf

[203] J. A. Schmitz, J. M. Sherman, S. Hansen, S. J. Murray, S. Balkir, and M. W. Hoffman, “A Low-Power, Single-Chip Electronic Skin Interface for Prosthetic Applications,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 13, no. 6, pp. 1186–1200, dec 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8876666/

[204] S. Kim, H. Shin, K. Song, and Y. Cha, “Flexible piezoelectric sensor array for touch sensing of robot hand,” in *2019 16th International Conference on Ubiquitous Robots (UR)*. IEEE, jun 2019, pp. 21–25. [Online]. Available: https://ieeexplore.ieee.org/document/8768644/

[205] “Home | Tecnalia.” [Online]. Available: https://www.tecnalia.com/en/

[206] Y. Abbass, S. Dosen, L. Seminara, and M. Valle, “Full-hand electrotactile feedback using electronic skin and matrix electrodes for high-bandwidth human–machine interfacing,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 380, no. 2228, p. 23, jul

2022. [Online]. Available: https://royalsocietypublishing.org/doi/10.1098/rsta.2021.0017

[207] S. W. Smith, “Moving Average Filters,” in *Digital Signal Processing*. Elsevier, 2003, pp. 277–284. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/B9780750674447500522

[208] J. Malešević *et al.*, “Encoding contact size information with static and dynamic electrotactile stimulation delivered to the fingertip,” *Experimental Brain Research*, 2024.

[209] G. Gu, N. Zhang, H. Xu *et al.*, “A soft neuroprosthetic hand providing simultaneous myoelectric control and tactile feedback,” *Nature Biomedical Engineering*, 2021.

[210] B. Peerdeman, D. Boere, H. Witteveen, R. Huis in ‘tVeld, H. Hermens, S. Stramigioli, H. Rietman, P. Veltink, and S. Misra, “Myoelectric forearm prostheses: State of the art from a user-centered perspective,” *The Journal of Rehabilitation Research and Development*, vol. 48, no. 6, p. 719, 2011. [Online]. Available: http://www.rehab.research.va.gov/jour/11/486/pdf/peerdeman486.pdf

Smart Human Machine Interface Using Piezoelectric Sensors and Artificial Intelligence

Hussein Bassal