

Politecnico di Torino

Politecnico di Torino

Master's Degree Programme in Biomedical Engineering

Academic Year 2025/2026

Graduation Session: September 2026

Multichannel Transcutaneous Electrical Stimulation for Non Invasive Tactile Feedback in Upper Limb Prostheses

Supervisors:

Prof. Elisa Donati
Prof. Giacomo Valle
Prof. Danilo Demarchi

Candidate:

Federico Greppi

This thesis was carried out at the Institute of Neuroinformatics (INI), University of Zurich and ETH Zurich.

Abstract

Neuromorphic and biomimetic technologies offer a way to bridge artificial systems and the natural encoding of touch. This is relevant not only for prosthetic limbs, but also for wearable haptics, virtual reality, teleoperation, and robotic interfaces, where tactile information must be transformed into artificial feedback. In upper limb amputation, the loss of tactile feedback forces users to rely more strongly on vision and indirect cues, increasing cognitive load and reducing the natural integration of the artificial limb. Transcutaneous electrical nerve stimulation offers a non invasive approach for restoring or augmenting tactile feedback. In targeted TENS, electrodes are placed over anatomical regions where stimulation can activate peripheral nerve branches and evoke sensations projected to the hand. However, because the stimulus must pass through the skin and surrounding tissues, fibre level selectivity is limited and the resulting percept depends strongly on electrode placement and individual sensitivity.

This thesis develops and evaluates a multichannel targeted TENS platform for studying artificial tactile feedback in 10 healthy participants. The setup consisted of a single channel constant current stimulator, a multiplexer for addressing different electrode pairs, a microcontroller for real time stimulation control, and a digital to analog converter for amplitude modulation. The main goal was to study how rich tactile information, represented as pressure profiles distributed over the hand, can be compressed into a small number of stimulation channels by modulating amplitude and frequency in real time.

Before testing dynamic feedback strategies, the stimulation space was calibrated for each participant. Detection thresholds, maximum comfortable amplitudes, and projected hand fields were measured to define subject specific operating ranges. Psychophysical characterization showed that amplitude and frequency play different perceptual roles: amplitude mainly controlled perceived intensity and spatial recruitment, whereas frequency affected temporal quality and contributed to intensity in a nonlinear way.

The central part of the work compared different strategies for translating simulated tactile interactions into stimulation commands. Pressure profiles over multiple anatomical hand zones were reduced to three stimulation channels corresponding to radial, medial, and ulnar projected territories. The tested encoders ranged from simple linear mappings to biologically inspired strategies. The most promising approach was an encoder that, instead of modelling tactile afferents in full detail, preserved a key biological principle: separating sustained pressure information from transient contact events. This strategy produced the clearest and most interpretable stimulation patterns in the task without visual information. Overall, the results suggest that effective biomimetic tactile feedback should adapt biologically inspired encoding principles to the constraints of transcutaneous stimulation, limited spatial selectivity, and subject specific perception.

Contents

Abstract **2**

List of Acronyms **8**

1 Introduction **9**

1.1 Motivation 9

1.2 Objective 10

1.3 State of the art 11

1.3.1 Biological encoding of touch 12

1.3.2 Approaches to restoring tactile feedback 14

1.3.3 Experimental evidence for non-invasive electrical feedback 14

2 Hardware Setup **17**

2.1 Block diagram 18

2.1.1 Digitimer Instruments 19

2.1.2 Microcontroller and D/A Converter 24

2.1.3 Hardware constraints 27

2.1.4 Laptop Requirements and USB Configuration 28

2.2 Bench validation with dummy load and oscilloscope 28

3 Control Software **31**

3.1 Repository organization 31

3.2 Hardware abstraction layer 32

3.3 Arduino firmware layer 33

3.4 Experiment GUI layer 33

3.5 Safety measures 33

4 Methods **35**

4.1 Electrode Selection and Placement 35

4.2 Threshold estimation and projected sensation mapping 37

4.2.1 Aim 37

4.2.2 Experimental design 37

4.2.3 Minimum referred area detection threshold 39
4.2.4 Maximum comfortable threshold and projected field mapping 40
4.3 Amplitude versus frequency discrimination on an ACR scale . . . 41
4.3.1 Aim 41
4.3.2 Experimental design 41
4.3.3 Psychometric analysis and Weber fraction 43
4.4 Frequency characterization 44
4.4.1 Aim 44
4.4.2 Experimental design 44
4.5 Overlapping train motion discrimination 46
4.5.1 Aim 46
4.5.2 Experimental design 46
4.6 Modulation strategy identification and naturalness experiment . . 48
4.6.1 Aim 48
4.6.2 Experimental design 48
4.6.3 Simulated tactile scenarios 50
4.6.4 Tested encoding strategies 51
4.6.5 Linear amplitude and linear frequency strategies 52
4.6.6 TouchMime encoding strategy 56
4.6.7 SA/RA LIF encoding strategy 63
4.6.8 Subject specific scaling and load matching 68

5 Results and discussion **70**
5.1 Threshold estimation and projected sensation mapping 70
5.2 Amplitude versus frequency discrimination on an ACR scale . . . 74
5.2.1 Linear ACR analysis 74
5.2.2 Adapted logarithmic ACR scale 76
5.3 Frequency characterization 78
5.4 Overlapping train motion discrimination 81
5.5 Modulation strategy identification and naturalness 84

6 Conclusion **90**

6.1 Main contributions 90
6.2 Limitations 91
6.3 Future possible directions 92

List of Acronyms

Acronym	Meaning
ACR	Activation charge rate
API	Application programming interface
BNC	Bayonet Neill Concelman connector
COM	Common Object Model
DAC	Digital to analog converter
DLL	Dynamic Link Library
DS8R	Digitimer DS8R constant current stimulator
D188	Digitimer D188 remote electrode selector
GUI	Graphical user interface
I2C	Inter Integrated Circuit communication bus
LIF	Leaky integrate and fire
PC	Personal computer
pps	Pulses per second
RA	Rapidly adapting
SA	Slowly adapting
TENS	Transcutaneous electrical nerve stimulation
tTENS	Targeted transcutaneous electrical nerve stimulation
TTL	Transistor transistor logic
USB	Universal Serial Bus
VR	Virtual reality

1 Introduction

1.1 Motivation

Tactile feedback is a central component of dexterous interaction with the environment. During grasping and manipulation, touch provides information about contact timing, contact location, pressure, object geometry, and texture. When this sensory channel is missing or degraded, as in upper limb amputation, users must rely more strongly on vision and indirect cues from the prosthesis or from the environment. This can reduce manipulation precision, increase cognitive load, and limit the sense that an artificial limb is naturally integrated into the body. Similar limitations also appear in broader wearable haptic interfaces, such as virtual reality or teleoperation systems, where tactile information must be artificially delivered to the user rather than sensed directly by the biological hand.

Electrical stimulation of the somatosensory system is one approach to restoring or augmenting tactile information. Compared with implanted peripheral nerve or cortical interfaces, transcutaneous electrical nerve stimulation (TENS) is attractive because it is non invasive, relatively simple to deploy, and suitable for iterative testing in healthy participants. In targeted TENS (tTENS), stimulation electrodes are positioned over anatomical regions where current can activate peripheral nerve branches and evoke sensations that may be perceived away from the electrode location. This makes tTENS a useful experimental platform for studying artificial tactile feedback without requiring surgery. At the same time, its non invasive nature introduces important limitations: selectivity is lower than in implanted interfaces, current spread depends strongly on electrode geometry and tissue properties, and perception varies substantially across subjects and stimulation sites.

A useful sensory feedback system must evoke percepts that should be comfortable, interpretable, stable enough for repeated use, and informative about the external tactile event. This is challenging because natural touch is not represented by a single scalar variable; instead, it is encoded by distributed spatiotemporal activity across populations of tactile afferents with different sensitivities and adaptation properties. A simple linear mapping from pressure to one stimulation parameter can therefore provide functional information in some tasks, but it may not capture important components of tactile dynamics such as contact transients, sustained pressure, spatial recruitment, and motion like changes across the hand.

Biomimetic approaches reproduce selected principles of the biological sensory system, rather than mapping a sensor measurement directly to an output command. Neuromorphic approaches pursue a similar goal using computational or electronic elements that emulate neural processing, often representing information as discrete events or spikes instead of continuously sampled values. These approaches provide a useful framework for addressing this problem. Instead

of directly mapping a sensor value to stimulation amplitude or frequency, the tactile event can first be transformed into a neural like representation inspired by biological mechanoreceptors and afferent populations. Neuromorphic systems implement this transformation through event-driven sensing, in which information is generated mainly when the sensory input changes, and spiking computation, in which brief events represent and transmit neural activity. This can reduce redundant data transmission and power consumption while preserving the timing of tactile events. Compact tactile circuits have already reproduced sustained and transient receptor-like responses, illustrating a possible route from software encoding to embedded sensory hardware [20, 21, 22, 23]. This representation can then be compressed into stimulation commands that are compatible with the available hardware. Conceptually, the complete loop contains three blocks: tactile sensing and pressure or indentation profile input, computational encoding, and electrical stimulation. In the long term, this type of pipeline could support portable sensory neuroprostheses and wearable haptic systems.

This translation from tactile input to stimulation output cannot be designed independently of the interface used to deliver the stimulus. The hardware determines which stimulation variables can be changed in real time, how many channels can be addressed, how accurately events can be timed, and which safety limits must be respected. Moreover, tTENS perception is highly subject specific. For this reason, a calibration and characterization phase is required before comparing encoding strategies: it defines the operating space in which later pressure to stimulation mappings can be delivered safely and interpreted consistently.

1.2 Objective

The first objective of this thesis was to update and extend an existing tTENS hardware and software platform. A previous project had established a first control architecture based on a Digitimer DS8R stimulator and a Digitimer D188 electrode selector. The present work started from that foundation and updated the control layer to use the newer Digitimer Common Object Model (COM) application programming interfaces (APIs), replacing the older Dynamic Link Library (DLL) based approach. This allowed hardware configuration to be separated more clearly from experiment logic and made the platform easier to integrate with the graphical user interfaces used during the experiments.

The second objective was to characterize the capabilities and constraints of the multichannel stimulation hardware. Because a single DS8R stimulator is routed through a D188 multiplexer, the system does not provide independent simultaneous current sources. Multichannel stimulation is obtained by rapidly switching the selected output channel and triggering the stimulator in a time multiplexed sequence. The usable stimulation frequency is therefore constrained

by the timing of channel selection, amplitude preparation, triggering, and the number of active channels. Defining these constraints was necessary before using the platform for dynamic multichannel patterns.

The third and main objective was to develop and test a complete experimental pipeline on 10 healthy subjects. The goal was to study how rich tactile information, such as pressure or indentation profiles that could be acquired from sensors on a prosthetic hand, can be compressed into a very limited set of transcutaneous stimulation commands. In the present setup, the effective stimulation interface was reduced to three projected territories, radial, medial, and ulnar, and the main real time variables were stimulation amplitude and frequency.

Before comparing encoding strategies, the stimulation space of each participant had to be defined. The first part of the experimental pipeline therefore measured the minimum detectable and maximum comfortable stimulation amplitudes for each channel. It also recorded the projected fields, namely the regions of the hand in which participants perceived sensations even though the electrodes were placed at the wrist. The next experiments characterized how changes in amplitude and frequency affected perceived intensity and perceptual quality. A further experiment tested whether delayed stimulation across two channels could evoke apparent motion. Finally, these calibration and characterization steps were used to compare five encoding strategies that transformed pressure profiles generated from six simulated tactile interactions into commands for the stimulation hardware.

The novelty of this work lies in evaluating this complete pipeline on the same group of subjects. While previous studies often focus on isolated aspects of transcutaneous stimulation, such as mapping, threshold estimation, psychophysical characterization, or encoding, this thesis combines these steps into a single subject calibrated framework for upper limb non invasive tactile feedback.

In summary, this thesis combines hardware development, software implementation, subject specific psychophysical calibration, and evaluation of encoding strategies. Section 2, *Hardware Setup*, describes the hardware components and their operation. Section 3, *Control Software*, describes the control architecture that connects the experiment GUIs to the Digitimer devices and Arduino firmware. Section 4, *Methods*, describes the experimental protocols and the aim of each experiment. Section 5, *Results*, presents and discusses the results obtained from these experiments. Section 6, *Conclusion*, synthesizes the main contributions and identifies directions for future development.

1.3 State of the art

Recreating tactile feedback requires first understanding which information the intact somatosensory system extracts from contact with the environment. This section therefore begins with the biological encoding of touch and the peripheral

nerves that carry hand sensation. It then reviews the principal approaches used to restore tactile information and, finally, the experimental evidence for upper-limb tTENS in healthy participants and amputees.

1.3.1 Biological encoding of touch

Touch begins when specialized mechanoreceptive endings in the skin convert mechanical deformation, such as pressure, stretch, or vibration, into electrical activity. These receptors are connected to sensory neurons whose afferent fibres carry the resulting signals through peripheral nerves toward the spinal cord and brain, where the distributed activity is combined to support the perception of contact, force, shape, texture, and motion. The tactile afferents of the hand are classified along two independent dimensions [1, 2]:

- **Adaptation:** slowly adapting (SA) afferents maintain activity during sustained deformation, whereas rapidly adapting (RA) afferents respond mainly to changes such as contact onset, release, and vibration.
- **Receptive-field size:** type I afferents have small, well-localized fields, whereas type II afferents have larger fields with diffuse boundaries.

Combining these properties gives four principal afferent classes:

- **SA1 (Merkel):** sustained pressure, edges, curvature, and fine spatial form; small receptive fields and greatest sensitivity at low frequencies.
- **SA2 (Ruffini):** skin stretch and hand configuration; large receptive fields and sensitivity to distributed deformation.
- **RA1 (Meissner):** contact transients, motion, flutter, and incipient slip; small receptive fields and low-to-intermediate-frequency sensitivity.
- **RA2/PC (Pacinian):** small, fast vibrations transmitted through tissues and objects; very large receptive fields and peak sensitivity around 200–300 Hz [3].

Figure 1 relates these functional classes to their principal end organs and to their depth within glabrous skin. Merkel and Meissner endings are located superficially and support spatially localized responses, whereas Ruffini and Pacinian endings are deeper and integrate deformation over broader regions [2].

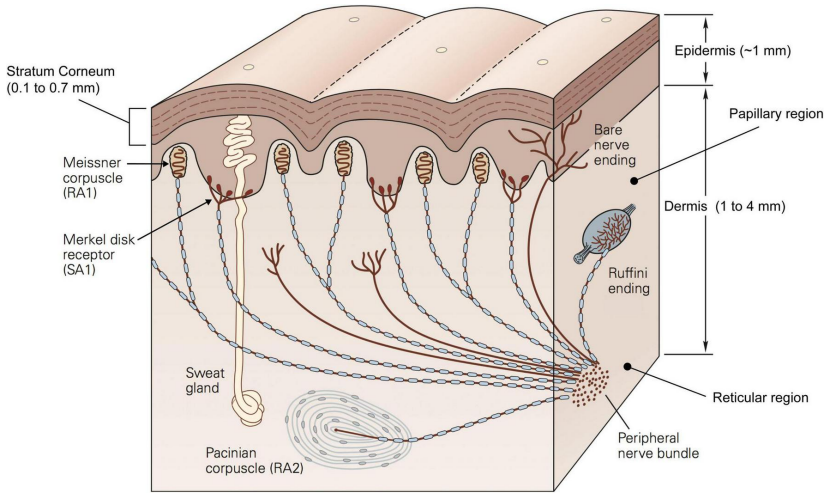


Figure 1: Organization of the principal cutaneous mechanoreceptors in glabrous skin and their associated afferent classes. Source: Deflorio et al. [2].

Sensory information from the hand is conveyed mainly by the median, ulnar, and radial nerves. Figure 2 shows their principal palmar and dorsal cutaneous territories. Stimulation of one of these afferent pathways can therefore evoke a sensation referred to the corresponding region of the intact or phantom hand. Although anatomical territories overlap and vary between individuals, this organization provides the basis for the radial, median, and ulnar stimulation sites used in this thesis and for the projected fields assessed during calibration [4].

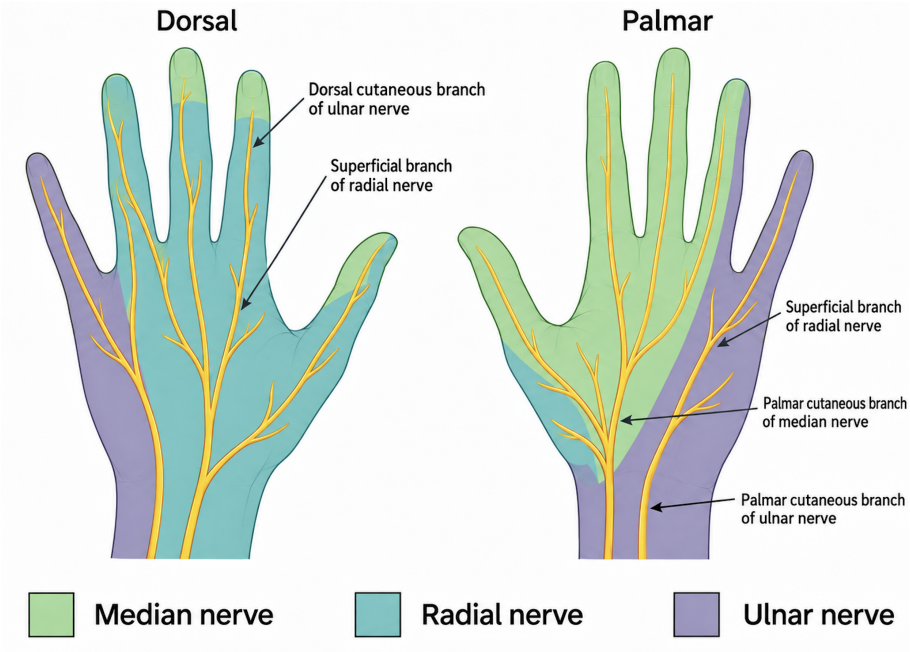


Figure 2: Principal cutaneous innervation territories of the dorsal (left) and palmar (right) surfaces of the hand.

1.3.2 Approaches to restoring tactile feedback

Artificial sensory feedback for an upper-limb prosthesis can be organized according to the site and method used to deliver information [5, 6]:

- **Stimulation of intact skin.** Mechanotactile interfaces apply pressure, indentation, or skin stretch and are intuitive but can be mechanically bulky. Vibrotactile interfaces are compact and effective for discrete events or force levels, although the sensation is normally localized at the actuator. Electrocutaneous interfaces activate cutaneous receptors or superficial afferents beneath skin electrodes. They allow thin multielectrode arrays and flexible modulation, but perception depends on placement and skin condition and is generally referred to the stimulated skin site.
- **Peripheral neural interfaces.** These interfaces act on peripheral nerves rather than on receptors at the feedback site. Non-invasive interfaces recruit nerve fibres through electrodes placed on the skin, whereas invasive interfaces place electrodes around, within, or close to a nerve. Implantation can improve selectivity and long-term access to somatotopic percepts, but requires surgery and clinical management [7, 8].
- **Reinnervative or central approaches.** Targeted sensory reinnervation surgically redirects residual nerves to accessible skin sites, while cortical interfaces stimulate somatosensory areas directly. These approaches may provide access when conventional peripheral stimulation is insufficient, but they are invasive and more clinically complex.

Within this classification, tTENS is a non-invasive peripheral neural interface. Electrocutaneous stimulation and tTENS both deliver current through surface electrodes, but they differ in their intended neural target: electrocutaneous feedback primarily activates local cutaneous afferents and is perceived near the electrode, whereas tTENS aims to recruit a proximal nerve pathway and evoke a referred sensation in its distal anatomical territory. Its main advantages are the absence of surgery, rapid electrode repositioning, compatibility with repeated psychophysical testing, and the possibility of producing somatotopic sensations in the intact or phantom hand. Its limitations include current spread, lower selectivity than implanted interfaces, sensitivity to electrode displacement and skin impedance, and substantial inter-participant variability in projected fields and comfortable stimulation ranges. Individual mapping and calibration are therefore essential [11, 12, 14].

1.3.3 Experimental evidence for non-invasive electrical feedback

The upper-limb tTENS literature includes studies in both healthy participants and amputees. Healthy cohorts are useful for characterizing stimulation parameters

and perceptual mechanisms under controlled conditions, whereas studies in amputees establish whether residual nerves can provide useful feedback from a prosthetic hand. The latter have demonstrated repeatable sensations referred to the phantom hand, localization to individual fingers or hand regions, graded representation of contact force, and temporal cues for events such as contact onset or slip. When incorporated into a closed loop, these sensations have supported force matching, recognition of object properties, and prosthesis manipulation [11, 12, 13, 14, 18, 19].

Table 2 summarizes representative studies and identifies the part of the feedback problem addressed by each one. MAP denotes the anatomical mapping of where an electrically evoked sensation is perceived; THR denotes measurement of perceptual threshold and usable stimulation range; JND denotes the smallest reliably discriminable change between two stimuli; FRQ denotes investigation of pulse-frequency effects; MCH denotes simultaneous or sequential use of multiple stimulation channels; DEL denotes manipulation of the delay between sites to create temporal or moving percepts; ENC denotes conversion of sensor information into stimulation commands; and APP denotes evaluation in a functional or prosthesis-control task. The experiments in this thesis address all of these categories, from calibration and discrimination to multichannel temporal patterns, encoding strategies, and application-oriented tactile scenarios.

Table 2: Representative studies of non-invasive electrical tactile feedback for upper-limb applications. Participant counts refer to the cohorts reported in each study.

Study	Participants	Stimulation target	Focus	Main contribution
Paredes et al. [9]	10 healthy	Forearm electrocutaneous	FRQ, JND, APP	Quantified frequency coding and its effect on closed-loop control.
Chai et al. [11]	11 amputees	Residual-limb sites	MAP, THR	Characterized location, quality, stability, and parameter dependence of evoked phantom-hand sensations.
D'Anna et al. [12]	4 amputees	Residual peripheral nerves	MAP, THR, ENC, APP	Integrated somatotopic TENS into a bidirectional prosthesis and modulated sensation intensity.
Osborn et al. [13]	2 amputees	Targeted residual nerves	MAP, THR	Demonstrated targeted tTENS for somatotopic phantom-limb feedback.

Study	Participants	Stimulation target	Focus	Main contribution
Shin et al. [14]	7 healthy; 1 amputee	Median/ulnar nerves at upper arm	MAP, JND, ENC	Elicited spatially distinct and graded tap/hold sensations in the hand or phantom hand.
Vargas et al. [15]	7 healthy	Proximal nerve grid	MCH, JND	Studied concurrent stimulation and merging of projected hand percepts.
Scarpelli et al. [16]	9 healthy	Multiple upper-limb sites	MCH, DEL	Quantified conditions for apparent moving sensations across the hand.
Hao et al. [17]	3 amputees	Stump projected-finger regions	MAP, ENC, APP	Related finger-specific stimulation to cortical activation and force-matching behaviour.
Vargas et al. [18]	8 healthy	Median/ulnar proximal nerves	MCH, ENC, APP	Used evoked feedback during prosthesis control for object-property recognition.
Scarpelli et al. [19]	5 amputees	Somatotopic TENS sites	ENC, APP	Encoded force and slip as complementary sustained and transient information.

The main contribution of this thesis is the evaluation of a complete tTENS pipeline in the same 10 participants. The protocol progresses from subject-specific calibration and psychophysical characterization to the comparison of multiple encoding strategies across different simulated tactile scenarios. This common cohort links the final strategy results directly to the perceptual constraints measured in the preceding experiments.

2 Hardware Setup

This chapter describes the physical platform used to deliver multichannel tTENS. The final setup combined a Windows computer running the experimental graphical interface, a Digitimer DS8R constant-current stimulator, a Digitimer D188 remote electrode selector, an Arduino/Genuino 101 microcontroller, and a 12-bit 0–10 V digital-to-analog converter (DAC). Device specifications and operating constraints were taken from the official DS8R and D188 operator's manuals [26, 27]. Figure 3 shows the laboratory configuration used during the experiments, including the Digitimer units, the control computer, the graphical interface, and the electrode connections on the participant's wrist.

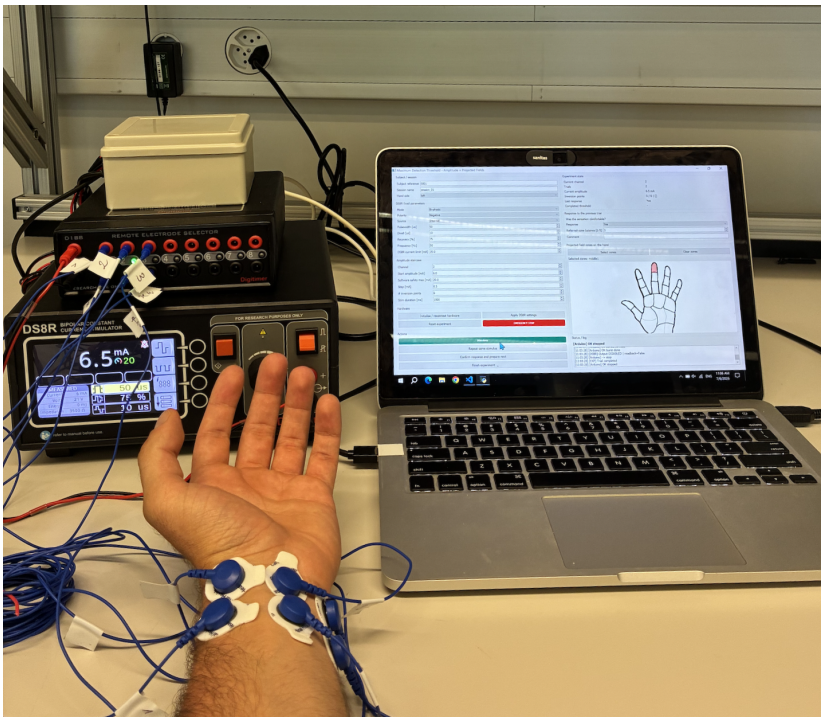


Figure 3: Photograph of the final laboratory setup used for multichannel tTENS experiments. The DS8R stimulator generates the current pulses, the D188 routes the output to the selected electrode pair, and the Windows computer runs the experimental interface and communicates with the hardware layer.

The platform was designed as a hierarchical stimulation system. The computer was responsible for high level configuration, graphical interaction, experimental flow, and data logging. The Arduino performed the operations that required deterministic timing: selecting the D188 channel, preparing the DAC output voltage, and triggering the DS8R. The DS8R generated the stimulation pulse, while the D188 routed that pulse to one of the available electrode pairs. Because the setup contains a single physical stimulator, multichannel stimulation is implemented by fast time multiplexing: at each stimulation event only one

electrode pair is connected to the DS8R output, and multichannel patterns are generated by switching across channels over time.

2.1 Block diagram

Figure 4 summarizes the final hardware architecture. The DS8R and D188 are connected to the computer through USB and controlled using the Digitimer COM APIs. These links are used to initialize the devices, configure the waveform parameters, select operating modes, set safety limits, and read back device states. However, the USB control path is too slow for real time updates of every stimulation pulse: DS8R software triggering is limited to about 10 pps, which is not sufficient for dynamic stimulation. For this reason, real time pulse scheduling is performed through the external hardware trigger and the Arduino.

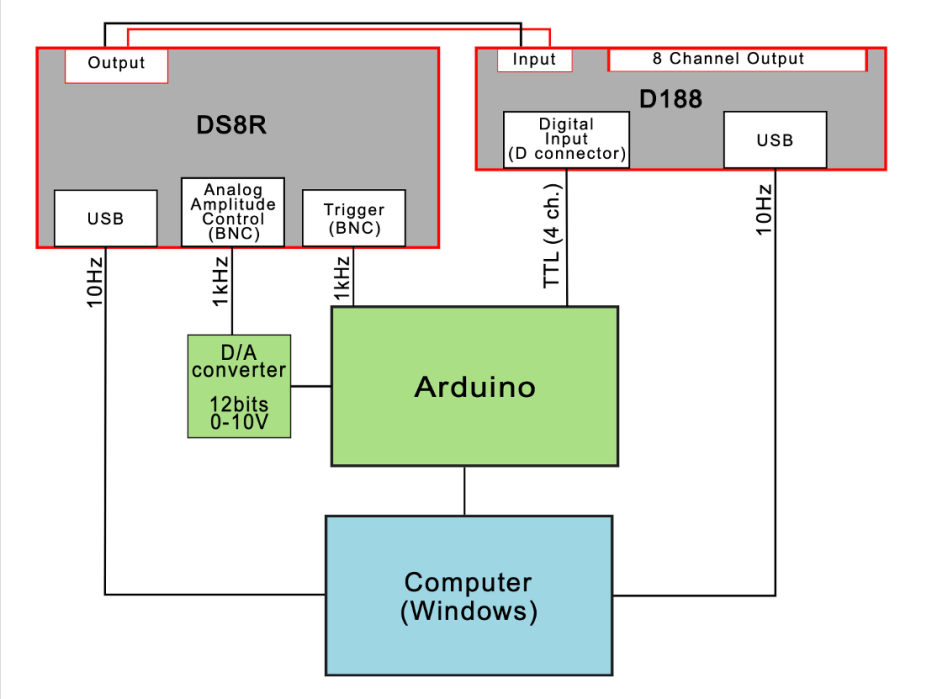


Figure 4: Overall block diagram of the time multiplexed tTENS platform. USB/COM communication is used for device configuration and readback, while the Arduino handles real time channel selection, DAC amplitude control, and DS8R hardware triggering.

The Arduino therefore acts as the real time executor. For each stimulation event, it selects the active D188 output channel through four TTL digital lines, sends the desired amplitude to the 12 bit DAC through I2C, waits for the control signals to settle, and sends a TTL trigger pulse to the DS8R. In this way, the Arduino defines which electrode pair is active, which current amplitude the DS8R should deliver, and when the pulse is generated. The 1 kHz value shown in the block diagram refers to the maximum hardware trigger rate accepted by the DS8R and

to the maximum update rate of the external analog-amplitude path. In practice, the usable rate of the complete multichannel system is lower, because each event also includes channel selection and DAC update.

A useful way to read the block diagram is to separate the stimulus path from the control path. The stimulus path is the high-voltage current path that starts at the DS8R output terminals, enters the common input of the D188, and exits through the selected D188 output channel toward the electrodes. The control path is instead composed of low-voltage signals used to configure and synchronize the hardware: USB/COM commands, serial commands to the Arduino, TTL address lines, the DS8R trigger input, I2C communication with the DAC, and the analog voltage sent to the DS8R amplitude control input. The DAC does not deliver stimulation to the participant; it only sets the DS8R analog amplitude control voltage. The actual stimulation current is always generated by the DS8R and only routed by the D188.

2.1.1 Digitimer Instruments

Two Digitimer instruments formed the core of the stimulation chain: the DS8R stimulator and the D188 remote electrode selector. Their functions were complementary. The DS8R defined the electrical pulse delivered to the participant, while the D188 defined which electrode pair received the pulse.

DS8R stimulator. The DS8R was used as the single constant-current stimulation source. It can generate monophasic or biphasic pulses with configurable current amplitude, pulse width, polarity, interphase interval, and recovery ratio. In the present work, the DS8R was configured in biphasic mode during stimulation protocols, while the variables that needed real time modulation were controlled externally: frequency through the trigger input and amplitude through the analog amplitude control.



Figure 5: Annotated DS8R front and rear panels.

The front panel in Figure 5 displays the parameters that define the stimulation waveform. The waveform-setting area (1) contains pulse width, recovery ratio, interphase interval, pulse mode, and polarity; the amplitude controls (2) define

the requested current and its source; the measured-pulse area (3) reports the most recent delivered values; and the configuration controls (4) provide access to the remaining device settings. The warning/status area (5) reports the output and safety state, while the output-enable control (6) arms or disables stimulation. During the experimental protocols, these parameters were not set manually from the front panel buttons; they were configured from the computer through the USB/COM interface. The display was still useful for visual inspection, because the operator could check that the selected waveform settings and the delivered pulse measurements were consistent with the protocol.

The DS8R amplitude can be controlled in two ways. For slow operation, the requested current can be set through the software interface. However, software updates and software triggering are limited to about 10 pps, which is too slow for the dynamic stimulation patterns used in this thesis. For real time amplitude modulation, the DS8R was therefore used in external analog control mode: the DAC provided a voltage between 0 and 10 V, and this voltage set the current amplitude of the next hardware triggered pulse. The measured pulse panel and warning indicators were continuously observed by the operator to verify correct stimulation delivery and detect possible safety warnings.

The rear panel in Figure 5 contains the connections used by the real time control architecture. The trigger BNC input (7) receives the TTL pulse generated by the Arduino and determines when the stimulus pulse is emitted. The analog control BNC input (8) receives the 0–10 V DAC signal and determines the amplitude of the next triggered pulse. The USB connection (9) is used for configuration and readback through the Digitimer COM API, but not for high-rate pulse scheduling.

Table 3: Main DS8R stimulation and control parameters relevant to the hardware platform.

Parameter	Value/Options	Information and use in this setup
Current output / Demand	0–1000 mA; 0.1 mA step; approximately ±2% accuracy.	Defines the requested current. During dynamic protocols, current was modulated through external analog control rather than repeated software demand updates.
Current amplitude limit	10–999 mA; 1 mA step.	Hardware safety ceiling used in every experiment to prevent unintended high-amplitude stimulation. Each experiment GUI also imposed an additional software ceiling, so that both device-level and software-level limits constrained the maximum deliverable current.
Pulse mode	Monophasic or biphasic.	Fixed to **biphasic** in all experiments to reduce net charge accumulation at the electrode–skin interface.
Polarity	Positive, negative, or alternating.	Fixed to **negative**; the negative stimulus phase therefore preceded the positive recovery phase.
Pulse width	50–2000 μs; 10 μs step; approximately ±2% accuracy.	Fixed to **50** μs for the stimulus phase in all experiments.
Interphase interval	1–990 μs; 10 μs step; approximately ±2% accuracy.	Fixed to **10** μs between the stimulus and recovery phases.
Recovery phase ratio	10–100%; 1% step; approximately ±2% accuracy.	Fixed to **75%**. The recovery amplitude was 75% of the stimulus amplitude and its duration was increased accordingly to approximately 66.7 μs to preserve charge balance.
Hardware trigger	TTL-compatible positive-edge input, +3 to +15 V; minimum pulse duration 5 μs; maximum hardware trigger rate 1000 pps.	Used by Arduino for precise pulse timing and frequency modulation.
Software trigger	Maximum software triggering rate approximately 10 pps.	Used only for manual or diagnostic triggering, not for real time stimulation scheduling.
External analog amplitude input	0–10 V input; linear scaling from 0 mA to the selected current amplitude limit; specified lag 1 ms; specified analog-control accuracy ±1 mA.	Driven by the DAC for pulse by pulse amplitude modulation. The ±1 mA term is an absolute accuracy specification; its percentage value depends on the selected current limit or on the commanded current used for normalization.
Compliance limit	400 V.	If the load impedance is too high, the DS8R may not be able to deliver the requested current.
Energy limit	300 mJ per pulse.	Protects against excessive pulse energy.

Figure 6 shows the standard pulse configuration used throughout the experiments. Pulse mode, polarity, pulse width, interphase interval, and recovery ratio remained constant, whereas stimulation amplitude and repetition frequency were varied according to the protocol and encoding strategy.

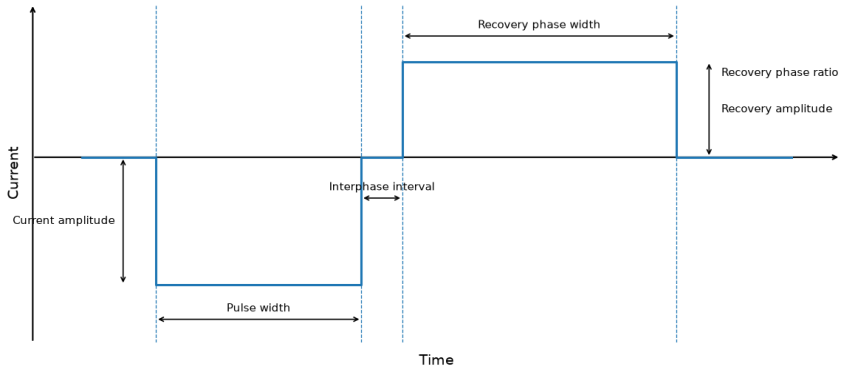


Figure 6: Biphasic DS8R waveform used in the experiments: a 50 μs negative stimulus phase, a 10 μs interphase interval, and a charge-balancing recovery phase at 75% of the stimulus amplitude.

The DS8R display and readback values are part of the safety and validation workflow. The operator can monitor the requested waveform parameters and the most recent pulse measurements, including current, voltage, impedance, and energy when available. The warning/status area (5) in Figure 5 is especially important because it indicates whether the stimulator is enabled, disabled, out of compliance, triggered too fast, or limited by pulse energy protection. These indicators are summarized in Table 4. If a critical warning appears, stimulation should be interrupted and the electrode contact, cables, and load should be checked before continuing.

Table 4: DS8R status indicators and warnings relevant during stimulation.

Icon	Indicator / warning	Meaning for the operator
	Output enabled	The output is enabled and a valid trigger will produce a stimulation pulse.
	Output disabled	The output is disabled, or the unit has disabled the output after a warning or fault. No pulse should be expected until the output is enabled again.
	Out of compliance	The load impedance is too high for the stimulator to deliver the requested current. Electrode contact, skin preparation, cables, and requested amplitude should be checked.
	Too fast	The DS8R is being triggered before the previous pulse waveform has completed. The trigger rate must be reduced or the waveform duration shortened.
	Pulse energy limit	The requested pulse would exceed the energy limit. Current amplitude, pulse width, or electrode/load impedance should be reduced before continuing.

D188 electrode selector. The D188 was used as an electrode selector: it does not generate stimulation but it routes the DS8R output to one selected output channel among eight possible electrode pairs. This is the element that enables multichannel stimulation while using only one stimulator.



Figure 7: Annotated D188 front and rear panels.

The front panel in Figure 7 shows the stimulation routing path. The DS8R output is connected to the common D188 input (1). The selected channel is then routed to one of the eight output electrode pairs (2). The rear panel contains the USB connection (3), which is used for configuration and readback, and the digital control connector (4), which receives the four TTL lines from the Arduino during real time stimulation.

Table 5: Main D188 parameters relevant to the time multiplexed setup.

Parameter	Value/Options	Information and use in this setup
Control modes	OFF / USB / 8TTL / 4TTL	USB was used for configuration and 4TTL for fast external channel addressing.
Logic input	TTL logic, 5 V	Arduino digital outputs drive the four address lines.
Maximum stimulus input voltage	400 V	Compatibility limit for the stimulation path routed from the DS8R.
Maximum stimulus input current	1 A	Compatibility limit for the stimulation path.
Minimum switching interval	< 1 ms; stimuli must not be applied during switching.	The firmware inserts a guard interval before triggering after channel selection.

The channel coding used by the Arduino to address the D188 outputs is reported in Table 6.

Table 6: D188 4TTL channel coding used for channel selection.

D188 output channel	4TTL code
CH1	1000
CH2	0100
CH3	1100
CH4	0010
CH5	1010
CH6	0110
CH7	1110
CH8	0001

A practical timing issue appeared during 4TTL channel selection. The Arduino/Genuino 101 firmware writes the four digital pins sequentially, not atomically (what does atomically mean? metteri together o qualcosa dle genere) through a single register operation. During a transition from one channel address to another, the D188 processes the digital input changes as they arrive. If the next bit changes before the processing of the previous transition has completed, the selector can miss one transition or briefly pass through an unintended intermediate address. Early tests showed wrong or unstable channel selections consistent with this behavior. The robust solution was to insert an inter-bit delay of about 10 μs between successive digital writes. This gives the hardware enough time to process one line transition before the next one is applied. This value is a validated implementation choice for the present microcontroller and wiring, not a universal D188 timing specification.

2.1.2 Microcontroller and D/A Converter

The Arduino/Genuino 101 was introduced to move deterministic, pulse by pulse operations away from the PC. Python and the COM interface are appropriate for setting parameters, organizing trials, and saving data, but they are not suitable for precise real time scheduling of channel switching and triggers. The Arduino therefore acts as the low level executor: it selects the D188 channel, sets the DAC amplitude, waits for the required guard intervals, and triggers the DS8R.

The wiring shown in Figure 8 implements the two control paths required for time multiplexed stimulation. In the annotated photograph, label 1 identifies the Arduino/Genuino 101 with terminal block shield, while label 2 identifies the GP8403 DAC module. Label 3 marks the four D188 TTL channel selection lines, connected to Arduino pins 10, 11, 12, and 13. Label 4 marks the DS8R trigger connection, driven by Arduino pin 8 together with the ground reference. Label 5 indicates the Arduino-side I2C and power terminals used for the DAC connection, including SDA, SCL, ground, and 3.3 V. Label 6 identifies the cable carrying these I2C and power lines to the DAC module. Finally, label 7 indicates the DAC output terminals used to provide the analog amplitude command to the DS8R external control input.

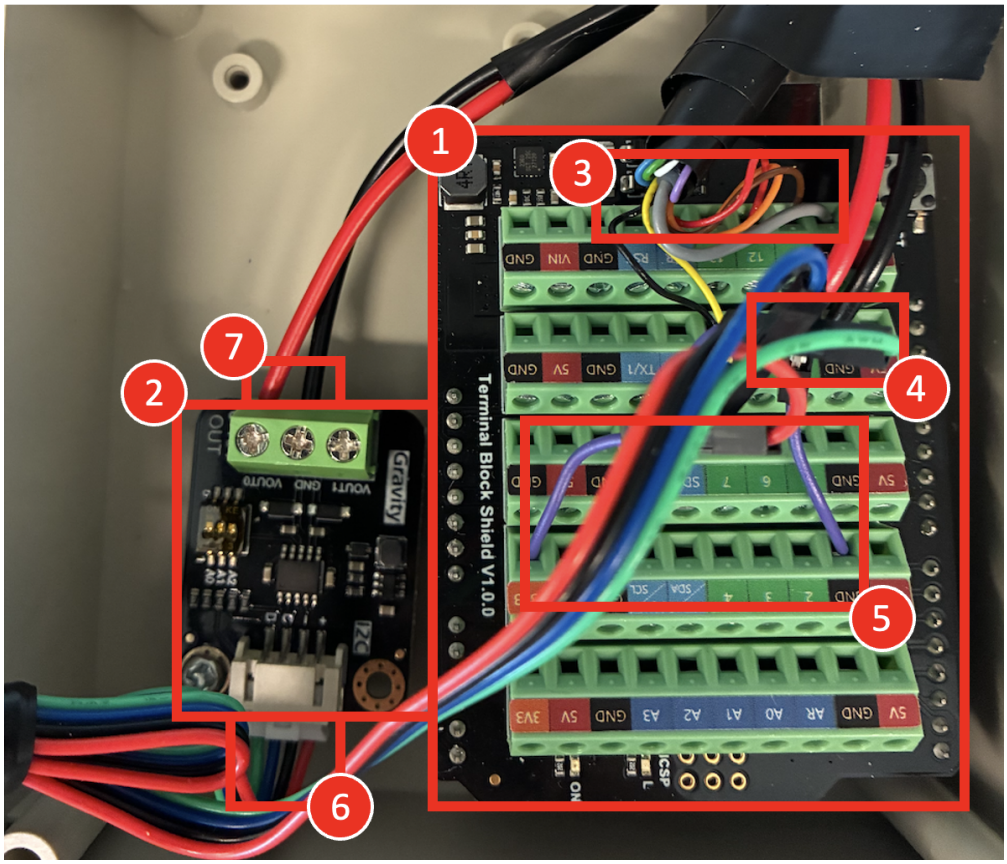


Figure 8: Arduino/Genuino 101 terminal shield and GP8403 DAC wiring used for real time stimulation control.

The first control path is digital: Arduino pins 10–13 drive the four D188 TTL address lines, allowing the firmware to select one of the D188 output channels before each stimulation event. Arduino pin 8 provides the hardware trigger sent to the DS8R, while the associated ground connection defines the low voltage reference for the trigger signal. The second path is analog: the GP8403 DAC is controlled through the I2C bus and receives its power and ground from the microcontroller side of the circuit. Its analog output is connected to the DS8R external amplitude control input, so the stimulation amplitude can be updated before each trigger without requiring a slow software command to the stimulator.

The external analog amplitude control input makes pulse by pulse amplitude modulation possible. In external source mode, the DAC voltage sets the requested DS8R output amplitude linearly between 0 V and 10 V. The upper end of this range corresponds to the current amplitude limit selected on the stimulator. If the limit is set to 50 mA, for example, 10 V corresponds to 50 mA rather than to the full 1000 mA range.

The 12-bit DAC resolution is sufficient for the intended amplitude modulation.

Over a 0–10 V range, the theoretical voltage resolution is

$$\Delta V = \frac{10\ \text{V}}{2^{12}} \approx 2.44\ \text{mV}.$$

If the DS8R current amplitude limit is set to 100 mA, this corresponds to an amplitude increment of approximately 0.024 mA per DAC step. This is below the DS8R nominal 0.1 mA current step and much smaller than the current output accuracy of the stimulator at most experimental amplitudes. Therefore, the DAC resolution is not the limiting factor for amplitude precision in the present setup.

One time multiplexed stimulation event consists of an ordered sequence of operations. First, the Arduino writes the D188 4TTL channel code using the validated inter-bit delay. Second, the firmware waits for the channel selection guard interval. Third, the DAC voltage corresponding to the requested current amplitude is prepared. Fourth, the analog path is allowed to settle. Finally, the Arduino sends the DS8R hardware trigger pulse. Channel specific amplitude modulation is possible only because the DAC value is updated before each trigger event, after the D188 channel has been selected.

Figure 9 illustrates one complete event using D188 channel 7 as an example. Starting from address 0000, TTL1, TTL2, and TTL3 are driven high sequentially, while TTL4 is written low, producing the final 4TTL address 1110. The 10 μs interval between successive writes prevents unstable address transitions. After the address has been completed, the firmware waits for the 100 μs channel-selection guard interval, updates the DAC command, allows 300 μs for analog settling, and retains a 100 μs timing margin before raising the DS8R trigger. The positive trigger edge initiates the biphasic stimulation pulse, while the trigger remains high for 1000 μs. The D188 manual specifies sub-millisecond switching and requires that stimulation is not delivered during a channel transition; the channel-selection guard interval was therefore included before triggering. The 10 μs inter-bit delay, 300 μs DAC-settling interval, 100 μs margin, and 1000 μs trigger-high duration are conservative design choices implemented in the firmware and verified on the oscilloscope, rather than timing values imposed by the instrument manuals. The horizontal spacing in the diagram is schematic rather than proportional to the physical durations.

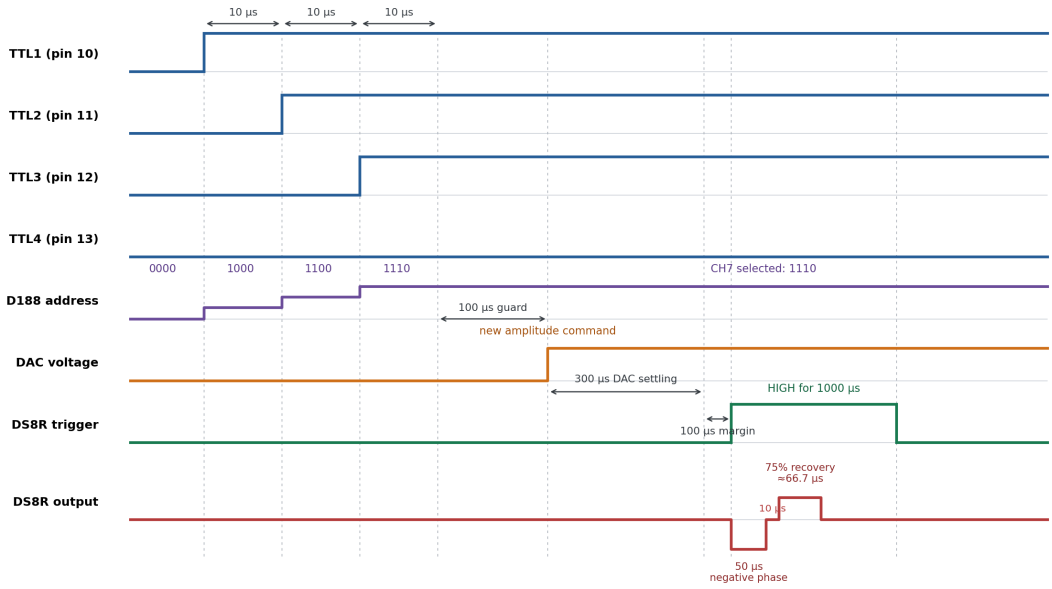


Figure 9: Timing diagram of one time-multiplexed stimulation event for D188 channel 7 (4TTL address 1110). Channel selection and the DAC amplitude update are completed before the positive edge of the DS8R hardware trigger.

2.1.3 Hardware constraints

In a multiplexed architecture, the limiting quantity is the total event rate, not the nominal frequency of one channel considered alone. If several channels are active, the scheduler creates one sequence of events distributed across channels. The total requested event rate is

$$f_{\text{total}} = \sum_{i=1}^{N_{\text{active}}} f_i,$$

where f_i is the requested event rate for channel i. A conservative feasibility condition is that the available event slot is longer than the time required to complete channel selection, amplitude update, trigger generation, and a timing margin:

$$T_{\text{slot}} = \frac{1}{f_{\text{total}}} > T_{\text{event}}.$$

For equally loaded channels, this gives

$$f_{\text{per channel,max}} \approx \frac{1}{N_{\text{active}} T_{\text{event}}}.$$

The conservative event time estimate used for safety checks is

$$T_{\text{event}} \approx 4T_{\text{bit}} + T_{\text{ch-settle}} + T_{\text{DAC-settle}} + T_{\text{trigger}} + T_{\text{margin}}.$$

The timing terms correspond to the sequence described in Figure 9: $T_{\text{bit}} = 10$ μs, $T_{\text{ch-settle}} = 100$ μs, $T_{\text{DAC-settle}} = 300$ μs, $T_{\text{trigger}} = 1000$ μs, and $T_{\text{margin}} = 100$ μs. Therefore,

$$T_{\text{event}} \approx 40 + 100 + 300 + 1000 + 100 = 1540 \ \mu\text{s},$$

which corresponds to an upper bound of about 649 events/s for the complete event stream. With three equally loaded active channels, this gives an approximate upper bound of 216 Hz per channel. The experimental GUIs therefore use a conservative maximum delivered frequency of 200 Hz, which remains below this firmware level estimate while leaving margin for serial communication, scheduler overhead, subject comfort, and device warnings. This value should be interpreted as an operating limit used in the implemented setup, not as the intrinsic maximum frequency of the DS8R alone.

2.1.4 Laptop Requirements and USB Configuration

The setup requires at least three USB connections: one for the DS8R, one for the D188, and one for the Arduino. Even when Arduino handles real time channel selection, amplitude preparation, and triggering, the Digitimer devices still need to be connected to the computer for configuration and mode selection.

USB reliability is a practical requirement for reproducibility. Communication with the DS8R, D188, and Arduino must remain stable during long experimental sessions, because failed configuration or delayed serial commands can interrupt the stimulation workflow.

The computer used to run the experimental software must be a Windows machine, because DS8R and D188 configuration relies on the Digitimer COM API.

2.2 Bench validation with dummy load and oscilloscope

Before applying the stimulation protocols on human participants, the stimulation chain was validated on a bench setup using a dummy resistive load. This test was performed for two purposes. First, it verified that the DS8R generated pulses with the expected amplitude and timing when operated without the D188. Second, it allowed the practical timing limitation introduced by the D188 multiplexing stage to be evaluated.

For the single-channel validation, the DS8R output was connected directly to a 1 kΩ dummy resistor, bypassing the D188. The stimulation parameters were set to the same waveform configuration used in the experimental protocols, namely biphasic stimulation with 50 μs pulse width, 75% recovery phase ratio, and 10 μs interphase interval. A stimulation amplitude of 1 mA and a frequency of 800 Hz were selected as a high-frequency test condition. With a 1 kΩ load, a 1 mA

current pulse corresponds to a voltage drop of approximately 1 V during the main phase.

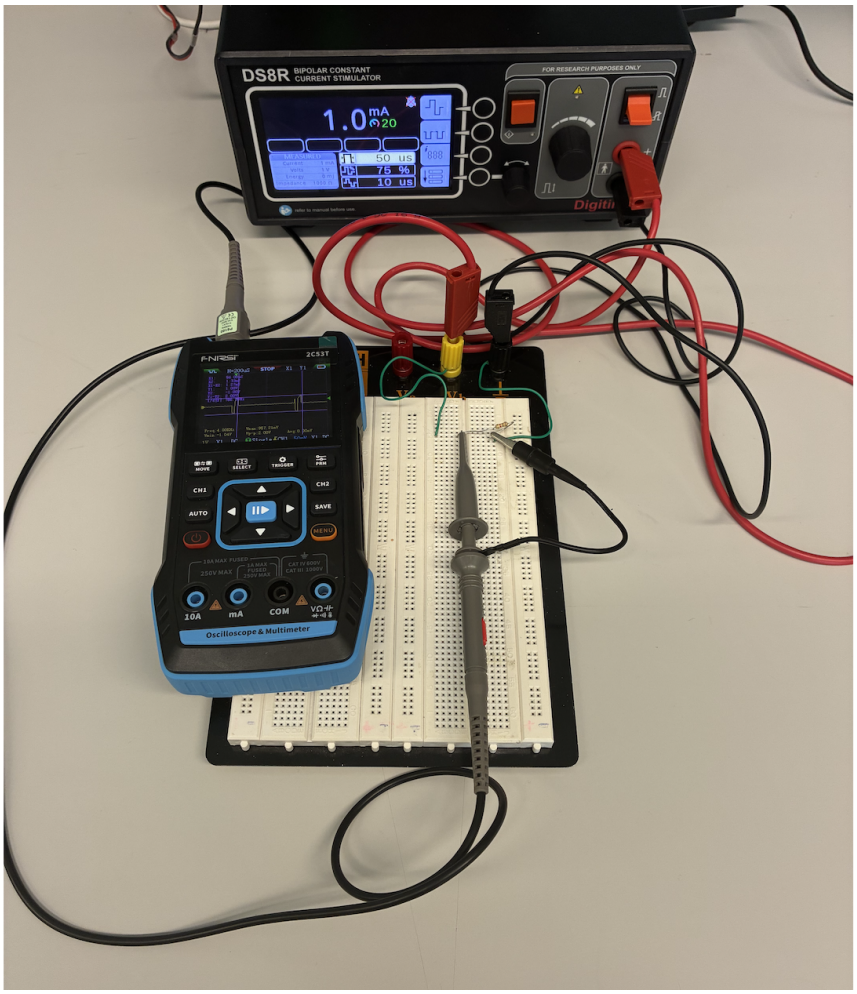
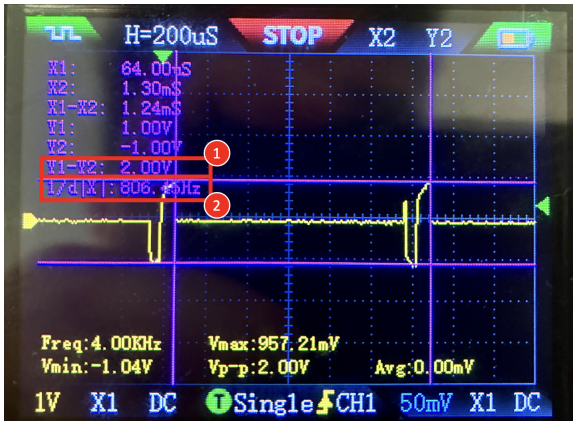


Figure 10: Bench setup used for single-channel validation. The DS8R output was connected directly to a 1 kΩ dummy resistor on the breadboard, and the voltage across the load was measured with the oscilloscope.

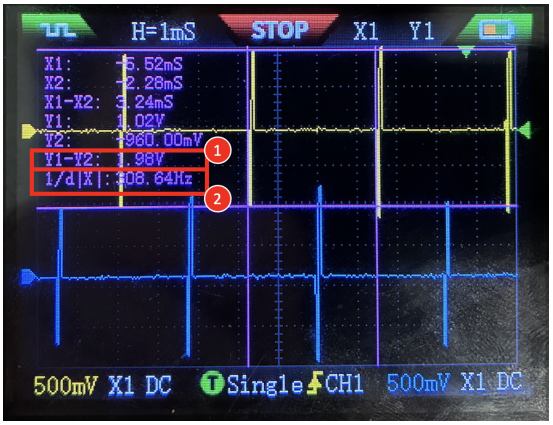
The oscilloscope measurement confirmed that the output voltage scale and the inter-pulse period were consistent with the selected stimulation parameters. In Figure 11a, annotation 1 indicates a peak-to-peak voltage of approximately 2 V across the dummy load, while annotation 2 indicates a measured frequency close to the commanded value of 800 Hz.

The multiplexed configuration was then tested by routing the DS8R output through the D188 and alternating stimulation between two output channels. In this condition, consecutive stimulation events were assigned to different channels, so the total event rate was twice the nominal per-channel rate. A nominal command of 400 Hz per channel was used as a stress test for the time multiplexed system. However, the oscilloscope measurement showed that the practical achievable per-channel rate was closer to 300 Hz, corresponding to a total event rate of

approximately 600 Hz. In Figure 11b, annotation 1 again indicates a peak-to-peak voltage of approximately 2 V, while annotation 2 indicates a measured per-channel frequency of about 308 Hz. This stress test was not performed with the full guard and safety intervals used in the final real time scheduler; it was intended to push the hardware chain close to its practical limit. The measured limit is therefore consistent with the theoretical timing budget and with the conservative frequency limit imposed in the experimental code, supporting the interpretation that the rate constraint is a real hardware and timing limitation of the multiplexed setup rather than an arbitrary software choice.



(a) Single-channel validation with the DS8R connected directly to a 1 kΩ dummy resistor. The measured inter-pulse interval confirms operation close to 800 Hz.



(b) Two-channel sequential stimulation through the D188. The measured timing indicates a practical per-channel rate of approximately 300 Hz, corresponding to a total event rate of about 600 Hz.

Figure 11: Oscilloscope validation of stimulation timing on a dummy load. The DS8R alone reproduced the expected high frequency pulse train, while the inclusion of the D188 introduced an additional timing constraint in the multiplexed condition.

3 Control Software

3.1 Repository organization

The software repository follows a modular organization. Reusable low level modules provide hardware communication, file handling, common widgets, and threshold loading, allowing the graphical interfaces and workflows of individual experiments to be implemented at a higher level without duplicating device-control code.

```
ttens/
|-- README.md
|
|-- API/
|   |-- d188_controller.py
|   |-- ds8r_controller.py
|   `-- __init__.py
|
|-- Arduino/
|   |-- firmware/
|   `-- archive/
|
|-- common/
|   |-- annotation_map.py
|   |-- file_utils.py
|   |-- hardware.py
|   |-- qt_widgets.py
|   `-- threshold_loader.py
|
|-- experiments_gui/
|   Experimental GUIs for calibration, mapping, frequency tests,
|   motion discrimination, and final scenario recognition.
|
|-- SIM/
|   Software simulations of the experimental hardware.
|
|-- datasheet/
|   Hardware manuals and technical documentation.
|
|-- pressure_profiles_23zones/
|   Synthetic pressure profiles for 23 hand zones.
|
|-- touchmime_outputs_13regions_fixed_capacity/
```

```
|   TouchMime firing-rate and recruitment profiles.
|
|-- recorded_data/
|   Automatically recorded experimental data.
|
|-- templates/
|   Anatomical images and GUI configuration resources.
|
|-- videos/
|   Unity videos for the tactile scenarios.
|
`-- test/
    Hardware and software diagnostic tests.
```

Separate resource folders contain the simulated pressure profiles, TouchMime outputs, anatomical templates, and Unity videos used to present the tactile scenarios in the final experiment. Experimental data, hardware documentation, diagnostic tests, and software simulations are likewise kept apart from the executable GUI and control modules.

The control software was organized in layers. The lowest layer consists of Python wrappers for the Digitimer COM APIs, which expose DS8R and D188 configuration commands in a compact and reusable form. Above this layer, Arduino serial commands are used to start, stop, and parameterize real time stimulation patterns. The highest layer consists of experiment specific PyQt graphical user interfaces, which manage trial logic, randomization, subject responses, safety buttons, and data logging.

3.2 Hardware abstraction layer

The script `API/ds8r_controller.py` wraps the DS8R COM API and provides methods to configure mode, polarity, amplitude source, current demand, current limit, pulse width, dwell, recovery, buzzer state, output enable, and software triggering. Similarly, `API/d188_controller.py` wraps the D188 COM API and exposes methods for selector mode configuration, front-panel indicators, delay/debounce settings, and USB-mode channel selection. In the final stimulation architecture, these APIs are mainly used for setup, verification, and safe device state management, while fast pulse timing is delegated to the Arduino firmware.

The shared hardware module in `common/hardware.py` combines these instrument control wrappers with Arduino serial communication. This module exposes higher level operations such as hardware initialization, safe shutdown, D188 channel selection through Arduino, DS8R amplitude update, DS8R output enable, and emergency stop. The experiment GUIs therefore do not need to duplicate fragile

device communication code.

3.3 Arduino firmware layer

The Arduino firmware implements the real time scheduler. It receives high level serial commands from Python, stores per-channel target frequencies and amplitudes, and executes the low level sequence required for each stimulation event: D188 channel selection, channel-settle delay, DAC update, DAC-settle delay, and DS8R trigger generation. Different frequencies on different channels are handled using a phase-accumulator logic: at each scheduler update, each active channel accumulates phase according to its target frequency, and a stimulation event is generated when the accumulated phase crosses the event threshold. This avoids blocking delays and allows channels with different requested rates to coexist within the same global event stream. The firmware also exposes safety-related commands such as `stop` and `off`. The `off` command stops active schedulers, disables triggering, sets the DAC output to zero, and clears the D188 channel selection.

3.4 Experiment GUI layer

The PyQt GUIs define the experimental workflow rather than the low level stimulation timing. Each experiment has a dedicated GUI script in `experiments_gui/`, so that mapping, threshold estimation, frequency characterization, motion discrimination, and final scenario recognition can each implement their own trial logic and response interface. However, all these scripts rely on the same shared modules for operations such as hardware initialization, comfort-range loading, parameter validation, subject/session handling, CSV saving, and safe shutdown. This layered organization avoids duplicating hardware control code inside every experiment and makes it easier to compare protocols using the same stimulation backend.

3.5 Safety measures

The hardware and software setup included several safety layers. The DS8R hardware limits defined the first protection level by constraining the maximum deliverable current and by reporting warnings related to load impedance or device state. On top of this, the graphical interfaces imposed software amplitude ceilings derived from the subject specific comfort range, so that requested amplitudes remained inside the calibrated operating region.

The operator also had access to Stop and Emergency Stop controls during the experiments. The standard stop action interrupted the ongoing stimulation

sequence, while the emergency path forced the system into a safe idle condition by stopping the Arduino scheduler, disabling further triggers, setting the DAC output to zero, and disabling or clearing the active D188 selection when required. During stimulation, the operator monitored the subject, kept the wrist posture stable, checked the skin when needed, avoided prolonged stimulation at the same site, and interrupted the protocol immediately in case of discomfort, irritation, unexpected sensation, or critical device warning.

4 Methods

4.1 Electrode Selection and Placement

Electrode selection and placement were critical aspects of the stimulation setup, because small changes in electrode geometry or position can substantially modify the evoked percept. Electrode size, spacing, orientation, skin preparation, and contact impedance influence current density, comfort, local under-electrode sensation, and the probability of obtaining a projected percept on the hand.

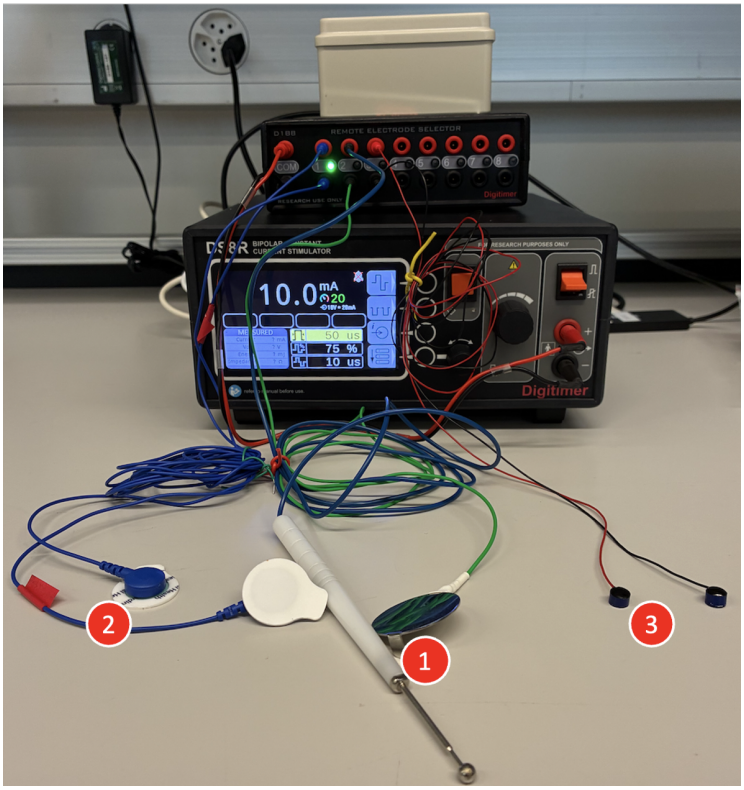


Figure 12: Electrode interfaces available during setup development.

The electrode interfaces considered during setup development are shown in Figure 12. The hand-held probe in Figure 12(1) consisted of a movable pen-like electrode with a rounded tip and a fixed circular return electrode with a radius of approximately 6 mm. Its small movable contact made stimulation more localized and allowed the wrist to be explored rapidly to identify sensitive sites that produced projected sensations in the hand. The probe was therefore used only for preliminary localization; the resulting candidate areas were then translated into standardized positions for the snap electrodes used in the experiments. The snap electrodes in Figure 12(2), with a radius of approximately 5 mm, and the sintered electrodes in Figure 12(3), with a radius of approximately 1–2 mm, were both tested during pilot sessions. Sintered electrodes provide a smaller and more

clearly defined stimulation surface, but require conductive gel and more careful skin preparation. For the final experiments, snap electrodes were selected because they reduced preparation time, improved comfort, and made repeated placement easier across subjects and sessions.

In the final experimental configuration, three stimulation channels were mainly used: radial, medial, and ulnar. Anatomically, these three wrist positions provided a compact way to address projected percepts distributed over the palmar hand. In particular, the radial and medial sites were intended to stimulate different regions of the median nerve, while the ulnar site was intended to target the ulnar nerve. This should not be interpreted as selective activation of individual fibres, but as a practical placement convention for transcutaneous stimulation. The actual projected fields were always verified experimentally, as described in the projected sensation mapping section.

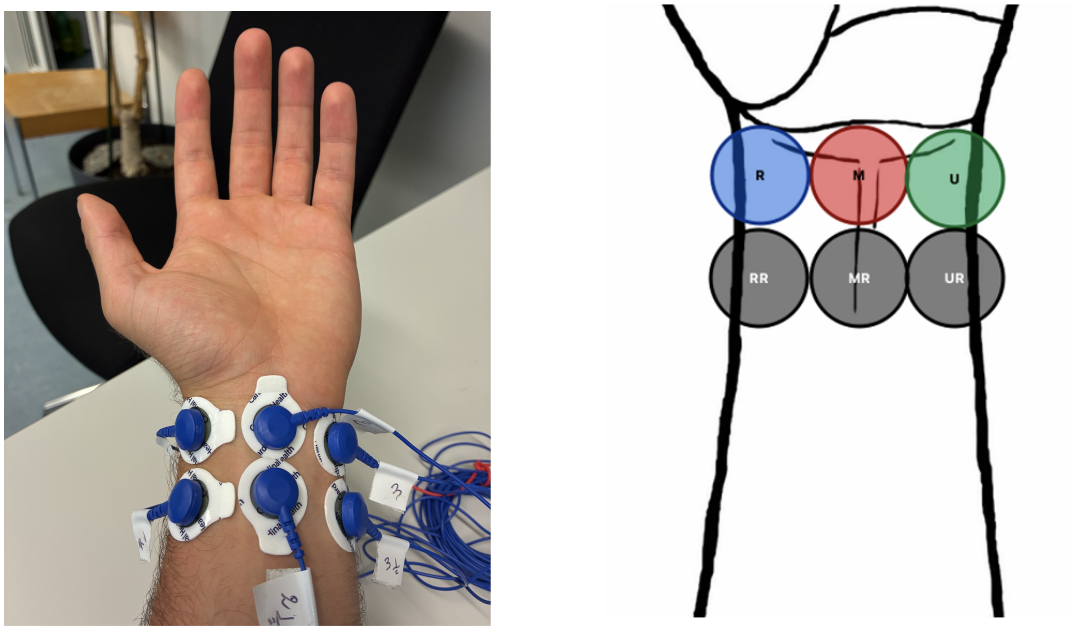


Figure 13: Final wrist electrode placement used for the three-channel stimulation configuration.

Figure 13 shows the placement convention used for the final three-channel configuration. The active stimulation sites were arranged along the radial, medial, and ulnar portions of the wrist, while the corresponding return electrodes were placed proximally. This arrangement was selected as a compromise between anatomical coverage of the hand, available space on the wrist, comfort during repeated stimulation, and the timing limits of the multiplexed hardware.

Each D188 output channel was treated as a channel-specific pair of stimulation electrodes connected to the two poles of the DS8R output. Because stimulation was delivered with a biphasic waveform, the recovery phase compensated the charge delivered during the stimulus phase, reducing net charge accumulation at

the tissue interface. In this configuration, the red and black DS8R outputs were therefore treated simply as the two poles of a bipolar stimulation output, rather than as fixed anodic or cathodic terminals.

The three-channel design also matched the timing constraints of the implemented stimulation scheduler. Oscilloscope validation showed a reliable total multiplexed throughput slightly above 600 Hz. Since the final protocols required stimulation frequencies up to hundreds of hertz, three active channels allowed a target of approximately 200 Hz per channel, whereas using more simultaneous channels would have reduced the available per-channel frequency margin.

4.2 Threshold estimation and projected sensation mapping

4.2.1 Aim

The first psychophysical experiment defined the subject specific operating range of each stimulation channel. This calibration was necessary because transcutaneous stimulation is strongly subject dependent: the same nominal current can produce different perceived intensities, projected hand fields, and local sensations under the electrodes across participants and electrode configurations. For this reason, each channel was characterized by a minimum referred area detection threshold and a maximum comfortable threshold. These two values defined the comfort range used for the following experiments:

$$CR_{ch} = T_{\max,ch} - T_{\min,ch}.$$

Here, $T_{\min,ch}$ is the lowest amplitude at which the subject reliably reported a sensation in the referred hand area, while $T_{\max,ch}$ is the highest amplitude that the subject still judged as comfortable.

4.2.2 Experimental design

The graphical interface used for the maximum comfortable threshold and projected field mapping protocol is shown in Figure 14. The left side of the interface contains the subject and session information, the fixed DS8R pulse parameters, the amplitude staircase settings, the hardware connection controls, and the main stimulation actions. The right side reports the current experimental state, stores the response to the previous trial, and displays the projected field zones selected on the hand map. This layout allowed the operator to control stimulation, monitor the staircase, and record both comfort and projected field information within the same interface.

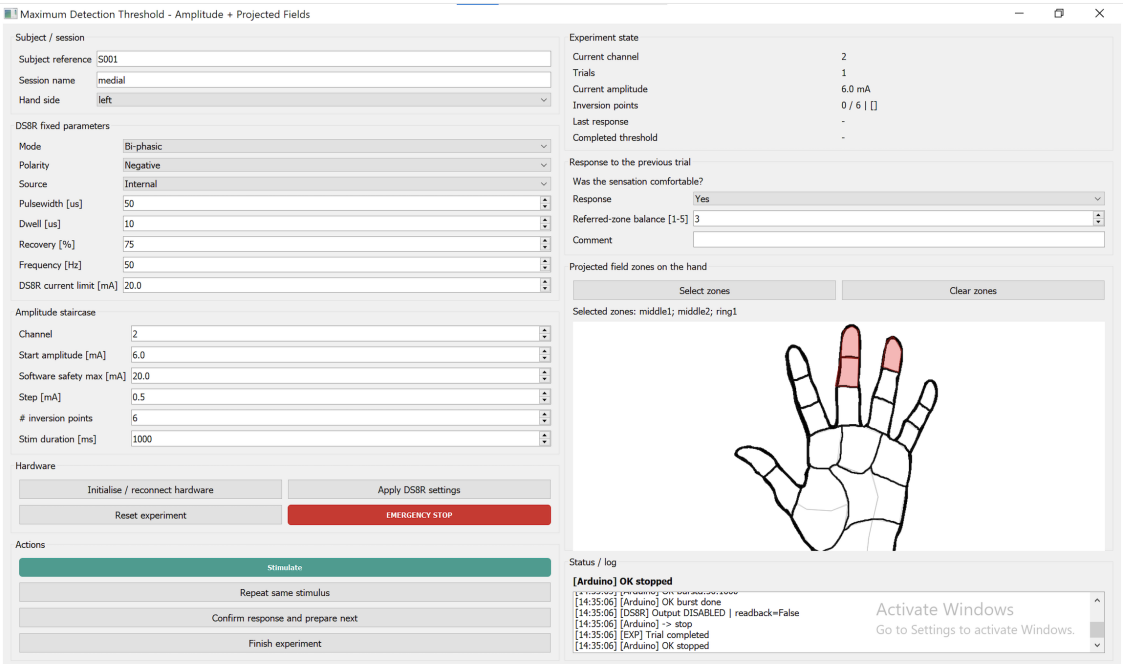


Figure 14: Graphical interface used for the maximum comfortable threshold and projected field mapping protocol.

Both thresholds were estimated using an amplitude staircase procedure. In a staircase, the stimulation amplitude is increased or decreased from one trial to the next according to the subject response, so that the sequence progressively approaches the perceptual boundary of interest. This method was selected because the goal was to identify a subject specific threshold rather than to sample the whole stimulation range uniformly. Compared with presenting random amplitudes, a staircase is more efficient for threshold estimation because the next stimulus is adapted to the previous response and the sequence rapidly converges toward the region where the response changes. A limitation is that the subject may partly infer the expected direction of the next stimulus, especially after repeated reversals, which can introduce a response bias or adaptation to the task structure. For this reason, the threshold was computed from inversion points rather than from the first reported response.

The stimulation waveform was kept fixed across trials, while only the current amplitude was updated according to the subject response. Stimulation frequency was fixed at 50 Hz, an intermediate value within the range subsequently explored in the frequency and encoding experiments. The fixed DS8R parameters were set in the corresponding panel of the GUI and are summarized in Table 7. An inversion was defined as a change in response between two consecutive staircase trials, and the final threshold estimate was computed as the mean of six inversion points. The resulting $T_{\max}$ should therefore be interpreted as a comfortable amplitude reference at 50 Hz, rather than as an absolute tolerability limit independent of frequency. A command combining $I = T_{\max}$ with a frequency above 50 Hz may

be perceived as more intense than the stimulus used to estimate $T_{\max}$ and must remain subject to the protocol-specific scaling and comfort checks.

Table 7: Main fixed parameters used during the threshold staircases.

Parameter	Value
Waveform mode	Biphasic
Polarity	Negative
Pulse width	50 μs
Dwell time	10 μs
Recovery	75%
Stimulation frequency	50 Hz
Burst duration	1000 ms
Amplitude step	0.5 mA
Threshold estimate	Mean of 6 inversion points

4.2.3 Minimum referred area detection threshold

The minimum threshold staircase estimated the lower bound of useful stimulation. The subject was asked whether the stimulus produced a sensation in the referred hand area. If the subject did not perceive a referred sensation, the amplitude was increased by 0.5 mA. If the subject reported a sensation in the referred area, the amplitude was decreased by 0.5 mA. The staircase continued until six inversion points were collected, and their mean was used as the detection threshold $T_{\min}$.

In pilot experiments, ten inversion points were initially collected and the final estimate was computed from the last six points, leaving the first reversals as an adaptation phase. This was reduced to six inversion points in the final protocol to shorten the calibration time. The change was justified by pilot data, where the inversion amplitudes were already relatively stable and the additional reversals provided mostly redundant information.

Importantly, the detection threshold was defined with respect to the referred sensation, not simply to any local sensation under the electrode. This distinction is relevant because, especially at low amplitudes, some subjects may perceive only a local tingling or pressure like sensation close to the electrodes. Such local sensations were not considered sufficient for the final sensory feedback application, where the goal is to evoke interpretable projected sensations on the hand.

Figure 15 shows a representative minimum threshold staircase. The current initially increases until the subject starts reporting a referred hand sensation. After the first changes in response, the staircase oscillates around the perceptual boundary, and the final threshold is obtained from the inversion amplitudes.

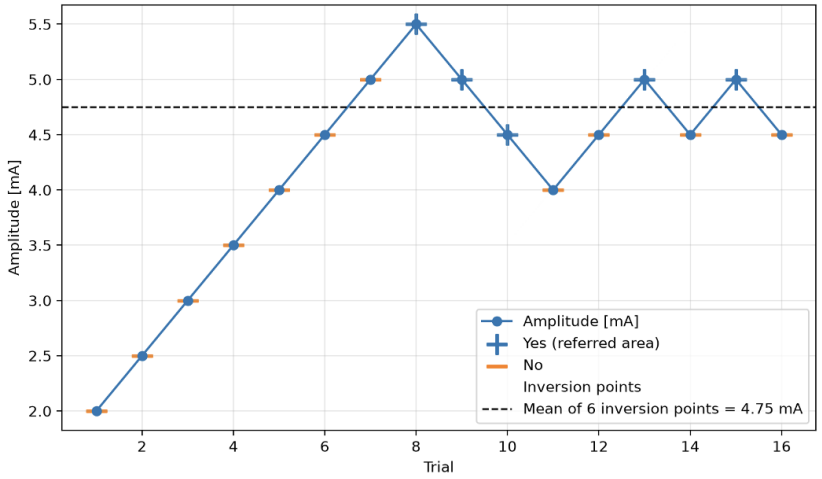


Figure 15: Representative minimum referred area detection threshold staircase. Plus symbols indicate trials in which the subject reported a referred sensation, while horizontal markers indicate trials without referred sensation.

4.2.4 Maximum comfortable threshold and projected field mapping

The maximum comfortable threshold was estimated with a similar staircase procedure, but the response question was different. After each stimulation burst, the subject was asked whether the sensation was still comfortable. This response was entered in the comfort response field shown in Figure 14. If the answer was yes, the amplitude was increased by 0.5 mA. If the answer was no, the amplitude was decreased by 0.5 mA. As for the minimum threshold, the final estimate $T_{\max}$ was computed as the mean of six inversion points.

During the maximum threshold staircase, the subject also reported the projected hand zones activated by the stimulation. This was done using the standardized hand map embedded in the GUI, where the operator selected the perceived zones after each trial. In addition, the subject rated the balance between local and referred sensation on a 1 to 5 scale. A score of 1 indicated that the sensation was perceived only under the electrode. A score of 3 indicated that the sensation was perceived equally under the electrode and in the referred hand area. A score of 5 indicated that the sensation was perceived only in the referred hand area. The selected zones and the balance score were saved together with the amplitude and comfort response for each trial.

Combining the maximum comfortable threshold staircase with projected field mapping allowed the experiment to capture not only the upper comfortable amplitude, but also how the perceived hand area evolved with current. This is important because increasing amplitude does not only increase perceived intensity; it can also recruit additional projected zones. Therefore, the comfort range should be interpreted not only as a scalar interval between $T_{\min}$ and $T_{\max}$, but also as a

range over which the spatial extent of the perceived sensation can change.

The results of this calibration experiment are reported in Section 5.1.

4.3 Amplitude versus frequency discrimination on an ACR scale

4.3.1 Aim

After subject specific calibration, and before testing the final encoding strategies, the role of the two stimulation variables modulated in real time was characterized: current amplitude and pulse frequency. The aim of this experiment was to measure how well subjects could discriminate changes in perceived intensity when either amplitude or frequency was varied. To make the two modulation variables comparable, the stimuli were expressed on an activation charge rate (ACR) scale. The charge-rate formulation was introduced by Graczyk et al. for peripheral nerve stimulation and was later adapted to transcutaneous feedback by Bucciarelli et al. [28, 29]. In the present work, the goal was to estimate the just noticeable difference (JND), defined as the difference between a reference stimulus and a test stimulus that is required for the subject to reliably perceive one of the two as stronger. By expressing both amplitude based and frequency based changes on the same ACR scale, the experiment aimed to determine which modulation variable provided finer perceived intensity discrimination.

4.3.2 Experimental design

The interface for the amplitude modulation condition is shown in Figure 16. The subject threshold section loaded the previously measured $I_{\min}$ and $I_{\max}$ values for the participant and stimulation channel under test and computed the reference amplitude and reference ACR. The ACR psychometric settings panel defined the tested ACR factors, the number of repetitions, the burst duration, the inter burst interval, and the frequency limits used by the protocol. The experiment state panel displayed the current ACR level, trial count, stimulus order, and delivered ratio, while the response panel implemented the two alternative forced choice answer by asking which of the two bursts felt stronger. Unless stated otherwise, the experimental GUIs were visible only to the operator, so that the participant could not use the displayed trial state or stimulation parameters as cues. The maximum-threshold protocol and the final scenario experiment are described separately where participant-facing visual information was required.

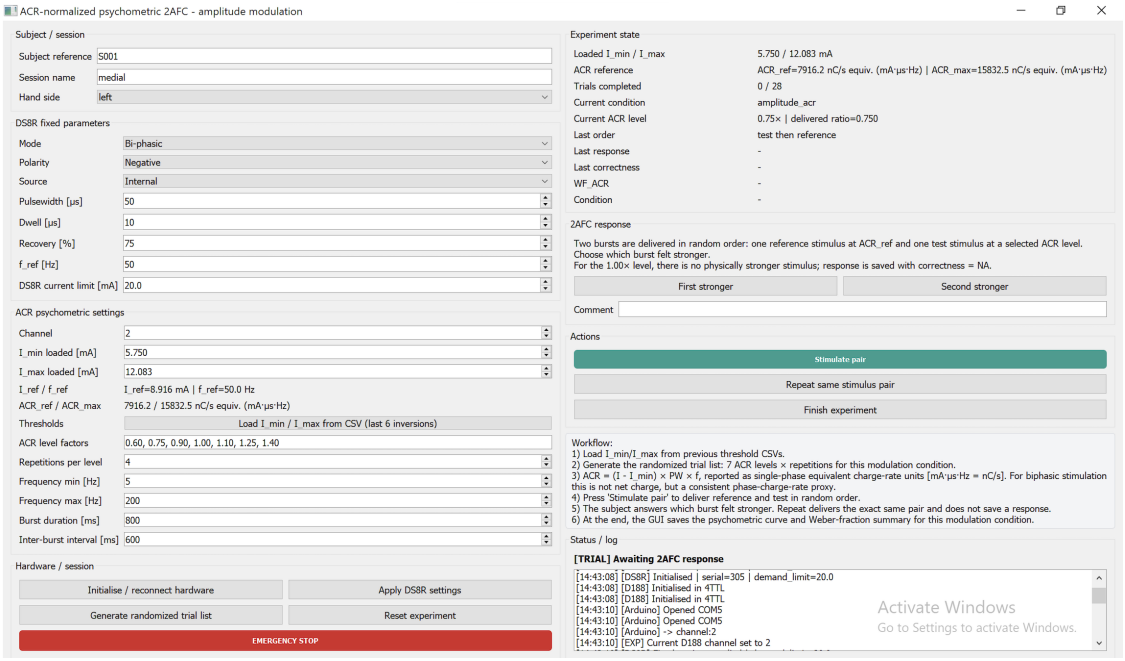


Figure 16: Graphical interface used for the ACR psychometric experiment in the amplitude modulation condition.

The experiment used a two alternative forced choice paradigm. In each trial, the subject received two stimulation bursts in random order: one reference stimulus and one test stimulus. The subject was always required to choose which burst felt stronger, even when the two stimuli had the same physical activation charge rate. The reference was defined at the middle of the subject specific comfort range:

$$I_{\text{ref}} = I_{\text{min}} + 0.5\,(I_{\text{max}} - I_{\text{min}})\,,$$

where I_{min} is the detection threshold and I_{max} is the maximum comfortable threshold. The reference frequency was fixed at $f_{\text{ref}} = 50$ Hz, and the pulse width PW was kept constant.

To compare the two modulation variables on a common scale, each stimulus was expressed in terms of activation charge rate:

$$\text{ACR} = (I - I_{\text{min}})\ \text{PW}\ f.$$

Here, $I - I_{\text{min}}$ represents the stimulation current above the detection threshold, PW is the pulse width, and f is the pulse frequency. Since $\text{mA} \cdot \mu\text{s}$ is equivalent to nC, multiplying by frequency gives units equivalent to nC/s. In this experiment, ACR was used as the charge delivered per second above the subject specific detection threshold. It is therefore not the total physical charge delivered by the stimulator, because the threshold current is subtracted and the biphasic waveform is charge balanced. However, it provides a useful normalized scale for comparing amplitude and frequency changes around the same operating point.

The reference ACR was therefore:

$$\mathrm{ACR}_{\mathrm{ref}} = (I_{\mathrm{ref}} - I_{\mathrm{min}})\ \mathrm{PW}\, f_{\mathrm{ref}}.$$

For each trial, the test stimulus was generated by multiplying this reference value by a factor k:

$$\mathrm{ACR}_{\mathrm{test}} = k\, \mathrm{ACR}_{\mathrm{ref}}.$$

Seven ACR factors were tested: $k = 0.60$, 0.75, 0.90, 1.00, 1.10, 1.25, and 1.40, as specified in the ACR psychometric settings panel of the GUI. Each level was repeated four times for each modulation condition, resulting in 28 trials each for amplitude and frequency modulation. The order of the trials and the order of the reference and test bursts were randomized. At $k = 1.00$, the reference and test stimuli had the same ACR; nevertheless, because the task was forced choice, the subject still selected one of the two bursts, and the response was saved without assigning correctness.

The two modulation conditions differed only in which stimulation parameter was adjusted to obtain the requested $\mathrm{ACR}_{\mathrm{test}}$. In the amplitude modulation condition, frequency was fixed at f_{ref}, while the test current was computed as:

$$I_{\mathrm{test}} = I_{\mathrm{min}} + \frac{\mathrm{ACR}_{\mathrm{test}}}{\mathrm{PW}\, f_{\mathrm{ref}}}.$$

In the frequency modulation condition, current was fixed at I_{ref}, while the test frequency was computed as:

$$f_{\mathrm{test}} = \frac{\mathrm{ACR}_{\mathrm{test}}}{(I_{\mathrm{ref}} - I_{\mathrm{min}})\,\mathrm{PW}}.$$

The frequency value was then rounded to the nearest deliverable integer frequency. With this design, the two conditions tested changes around the same reference ACR, but one implemented the change through current amplitude and the other through pulse frequency.

4.3.3 Psychometric analysis and Weber fraction

For each subject and modulation condition, responses were grouped by ACR ratio. The dependent variable was the probability that the test stimulus was judged stronger than the reference. The independent variable was the ACR variation from the reference:

$$x = 100 \left(\frac{\mathrm{ACR}_{\mathrm{test}}}{\mathrm{ACR}_{\mathrm{ref}}} - 1 \right).$$

For each subject, the empirical response probabilities were fitted with a logistic psychometric function:

$$P_{\mathrm{stronger}}(x) = \frac{1}{1 + \exp\left[-\frac{x - x_{50}}{s}\right]},$$

where x is the normalized ACR variation, x_{50} is the point at which the test stimulus is judged stronger with probability 0.5, and s controls the slope of the psychometric transition. A smaller value of s corresponds to a steeper curve and therefore to finer discrimination. The 75% point was obtained from the fitted function as the stimulus level at which $P_{\text{stronger}} = 0.75$. The JND was then defined as the difference between the 75% and 50% points:

$$\text{JND}_{\text{ACR}} = \text{ACR}_{75} - \text{ACR}_{50}.$$

The corresponding Weber fraction was computed as:

$$\text{WF}_{\text{ACR}} = \frac{\text{JND}_{\text{ACR}}}{\text{ACR}_{50}}.$$

The Weber fraction normalizes the JND by the reference operating point. Therefore, it expresses the minimum discriminable change as a relative fraction of the baseline stimulation level. A lower Weber fraction indicates finer perceptual resolution, because a smaller relative change in ACR is sufficient for the subject to judge the test stimulus as stronger. At group level, curves were obtained by fitting each subject separately and then averaging the fitted probabilities across subjects. Representative fitted psychometric curves and the group level outcomes are reported in Section 5.2, rather than in the methods section, to keep protocol description and results interpretation separated.

The results of the linear and adapted ACR analyses are reported in Section 5.2.

4.4 Frequency characterization

4.4.1 Aim

After the ACR and adapted ACRT psychometric analyses, an additional frequency characterization experiment was performed to describe how the quality of the evoked percept changes as a function of pulse frequency. The aim was to characterize how changes in pulse frequency affect not only perceived intensity, but also the temporal quality of the sensation.

4.4.2 Experimental design

The graphical interface used for this protocol is shown in Figure 17. The left side of the GUI contains the subject and site information, the fixed DS8R waveform parameters, the selected stimulation amplitude, and the list of frequencies and repetitions used to generate the randomized trial sequence. The right side displays the current experimental state, the response fields used to classify the perceived sensation and intensity, the stimulation actions, the workflow reminder, and the hardware log.

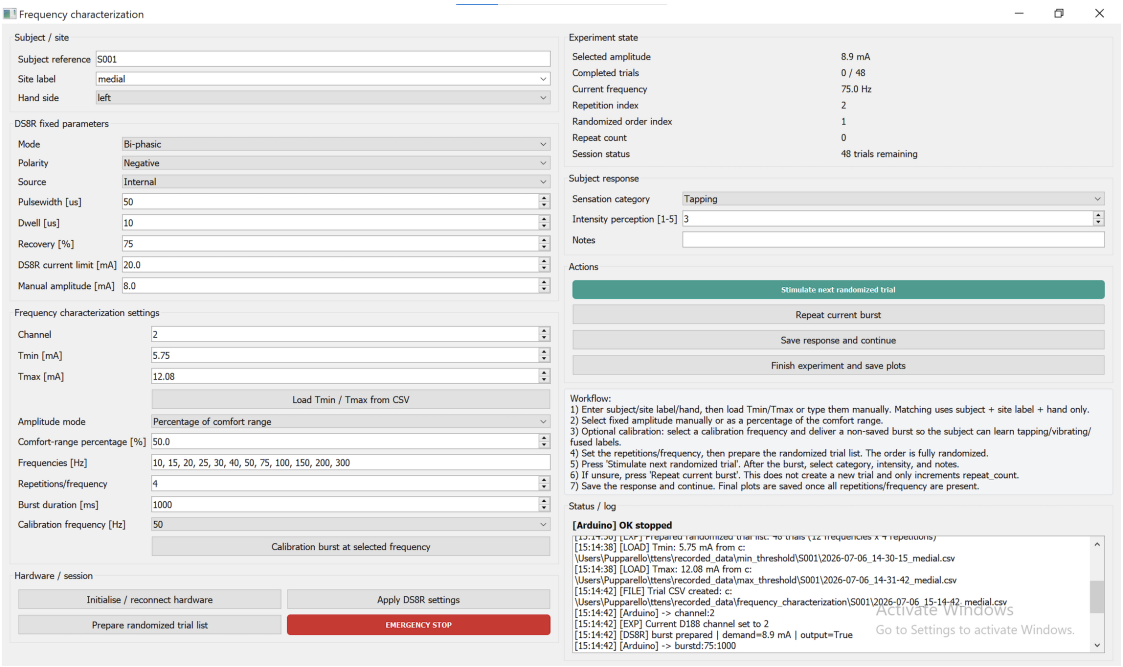


Figure 17: Graphical interface used for the frequency characterization experiment.

For each subject, the experiment was performed on one selected stimulation channel, the medial channel. The stimulation amplitude was kept fixed at 50% of the subject specific comfort range,

$$I_{\text{stim}} = I_{\text{min}} + 0.5\,(I_{\text{max}} - I_{\text{min}})\,,$$

where I_{min} is the detection threshold and I_{max} is the maximum comfortable threshold for the selected channel. Pulse width and waveform parameters were kept fixed, so that only the pulse frequency changed across trials. The tested frequencies were

$$f \in \{10, 15, 20, 25, 30, 40, 50, 75, 100, 150, 200, 300\}\ \text{Hz}.$$

Stimuli were delivered in randomized order. After each stimulus, the subject classified the sensation into one of three perceptual categories: *tapping*, corresponding to an intermittent or pulsatile sensation in which individual pulses could still be distinguished; *vibrating*, corresponding to a vibratory sensation in which individual pulses were no longer clearly resolved; and *fused*, corresponding to a faster and more merged percept. In addition to the category response, subjects rated the perceived intensity on a 1–5 scale.

The category responses were converted into probabilities. For each subject and frequency, the probability of each category was computed as

$$P_c(f) = \frac{N_c(f)}{N(f)},$$

where $N_c(f)$ is the number of trials classified as category c at frequency f, and $N(f)$ is the total number of valid trials at that frequency. Group level curves were obtained by first computing one probability curve per subject and then averaging across subjects.

Two transition frequencies were then estimated for each subject. The first transition, from intermittent to non-intermittent perception, was defined as the frequency at which the combined probability of vibrating and fused perception reached 50%:

$$P_{\text{nonint}}(f) = P_{\text{vibrating}}(f) + P_{\text{fused}}(f) = 1 - P_{\text{tapping}}(f) = 0.5.$$

The second transition, from non-fused to fused perception, was defined as the frequency at which fused perception reached 50%:

$$P_{\text{nonfused}}(f) = P_{\text{tapping}}(f) + P_{\text{vibrating}}(f) = 1 - P_{\text{fused}}(f) = 0.5.$$

Equivalently, this second criterion corresponds to $P_{\text{fused}}(f) = 0.5$. When a transition occurred between two tested frequencies, the transition frequency was estimated by linear interpolation.

The results of the frequency-characterization experiment are reported in Section 5.3.

4.5 Overlapping train motion discrimination

4.5.1 Aim

A final exploratory experiment was performed to evaluate whether two delayed stimulation trains delivered on two different mapped channels could evoke a directional or motion like percept. The aim was to estimate the delay required for subjects to discriminate the temporal order of the two channel activations, and to identify the delay range that produced the clearest sensation of movement across the hand.

4.5.2 Experimental design

The graphical interface used for this protocol is shown in Figure 18. The left side of the GUI defines the subject information, the two stimulation channels, the loaded comfort ranges, the selected amplitudes, and the tested onset delays. The right side displays the randomized trial state, the true train order used internally, the subject response buttons, the motion like clarity rating, and the workflow and log panels used by the operator during acquisition.

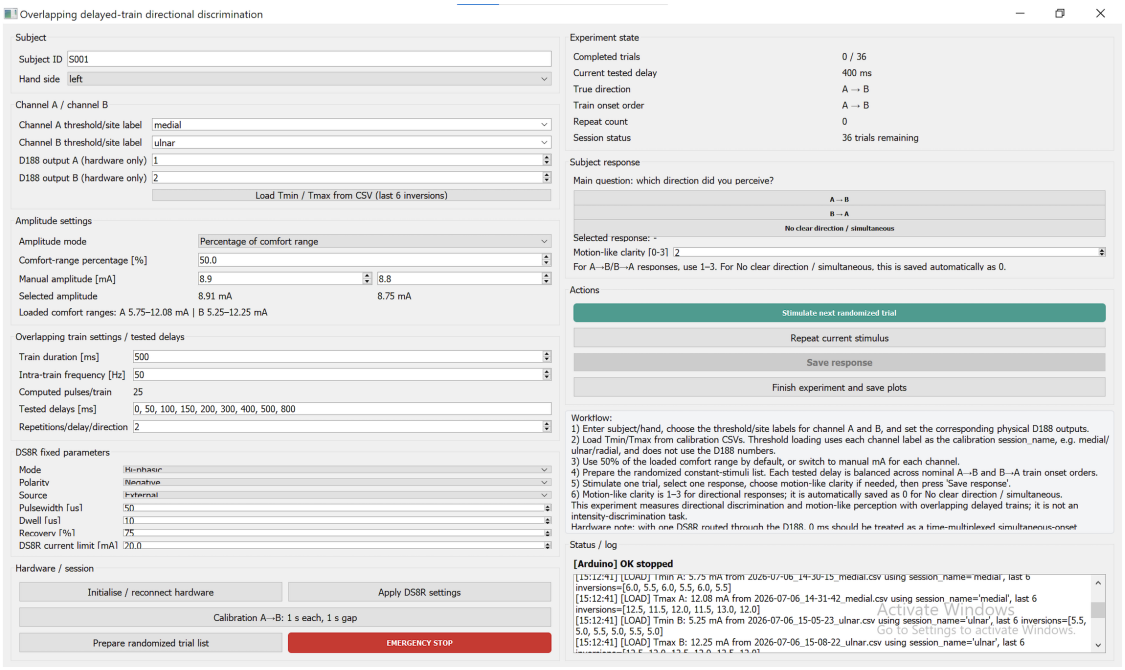


Figure 18: Graphical interface used for the overlapping delayed-train motion discrimination experiment.

Two stimulation channels were used: channel A corresponded to the medial site, and channel B corresponded to the ulnar site. Each trial consisted of two pulse trains with the same duration and frequency. One train started first, and the second train started after a delay d. The train duration was fixed at 500 ms and the intra-train frequency was fixed at 50 Hz. Stimulation amplitudes were set at 50% of the subject specific comfort range for each channel, unless manual adjustment was required.

The tested onset delays were

$$d \in \{0, 50, 100, 150, 200, 300, 400, 500, 800\}\,\text{ms}.$$

The 0 ms condition was included as a simultaneous onset control. Since the system used one DS8R stimulator routed through the D188 selector, this condition should be interpreted as a time multiplexed simultaneous onset approximation rather than true simultaneous stimulation from two independent outputs.

For each delay, both nominal directions were tested in randomized order. Each delay direction combination was repeated two times, giving

$$9\,\text{delays} \times 2\,\text{directions} \times 2\,\text{repetitions} = 36\,\text{trials}.$$

After each trial, the subject reported the perceived direction using three possible responses: A to B, B to A, or no clear direction. A separate motion like clarity rating from 0 to 3 was then recorded. This rating described how much the stimulus resembled a movement across the hand, rather than its intensity or

pleasantness. When the subject selected no clear direction/simultaneous, the motion like clarity rating was automatically saved as 0.

For each subject and delay, directional discrimination was quantified as the probability of a correct response:

$$P_{\text{correct},s}(d) = \frac{N_{\text{correct},s}(d)}{N_{\text{valid},s}(d)}.$$

For positive delays, a response was counted as correct when it matched the true train onset order. A descriptive temporal discrimination threshold was then computed as the smallest positive delay at which the subject reached at least 75% correct responses:

$$d_{75,s} = \min\left\{d > 0 : \ P_{\text{correct},s}(d) \geq 0.75\right\}.$$

Here, the 75% criterion represents a practical level above chance at which the subject could reliably identify which channel was activated first. When the criterion fell between two tested delays, the threshold was estimated by linear interpolation. Because the number of repetitions per delay was limited, $d_{75,s}$ should be interpreted as a descriptive index rather than a fitted psychometric threshold.

The results of the overlapping train motion-discrimination experiment are reported in Section 5.4.

4.6 Modulation strategy identification and naturalness experiment

4.6.1 Aim

This experiment was designed to compare different tactile encoding strategies during simulated hand object interaction scenarios. The goal was not only to evaluate whether the stimulation felt natural, but also whether the stimulation pattern was intuitively associated with the correct tactile event when visual information was removed.

4.6.2 Experimental design

Two graphical interfaces were used during this experiment. The main operator interface is shown in Figure 19. It contains the subject and block information, the selected tactile scenario, the chosen encoding strategy, the pressure reference and fixed frequency settings, the three channel mapping with the corresponding comfort ranges, and the hardware actions used to initialize, preview, randomize, and deliver the trials. The stimulation window is shown in Figure 20. During

visual feedback trials, this window displayed the Unity video of the simulated interaction while the encoded stimulation was delivered. During no visual feedback trials, the same window remained black, so that the subject received only the stimulation.

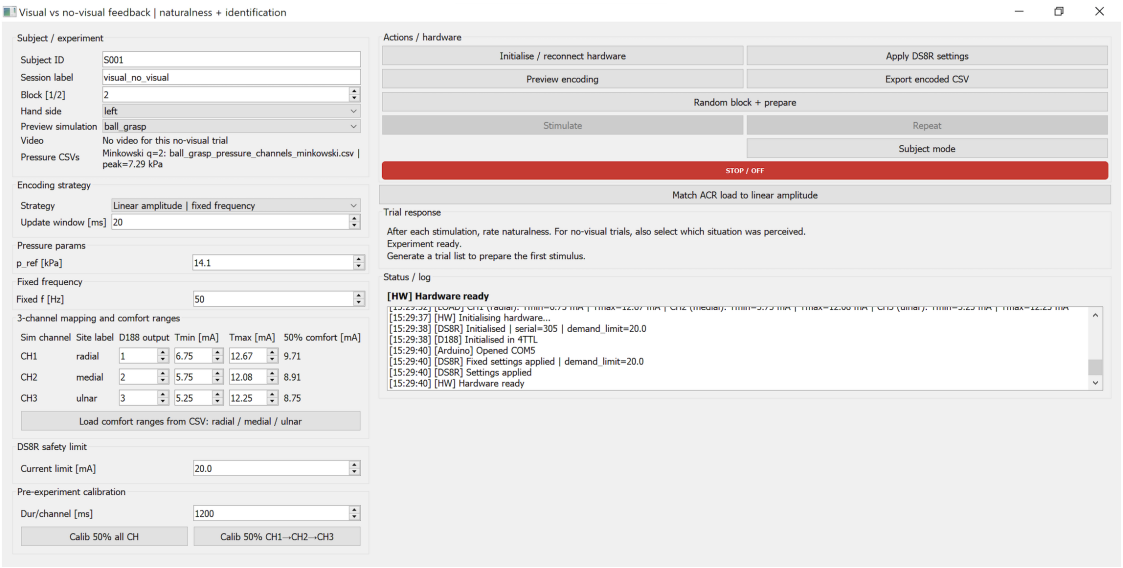


Figure 19: Main operator interface used for the visual versus no visual feedback experiment.

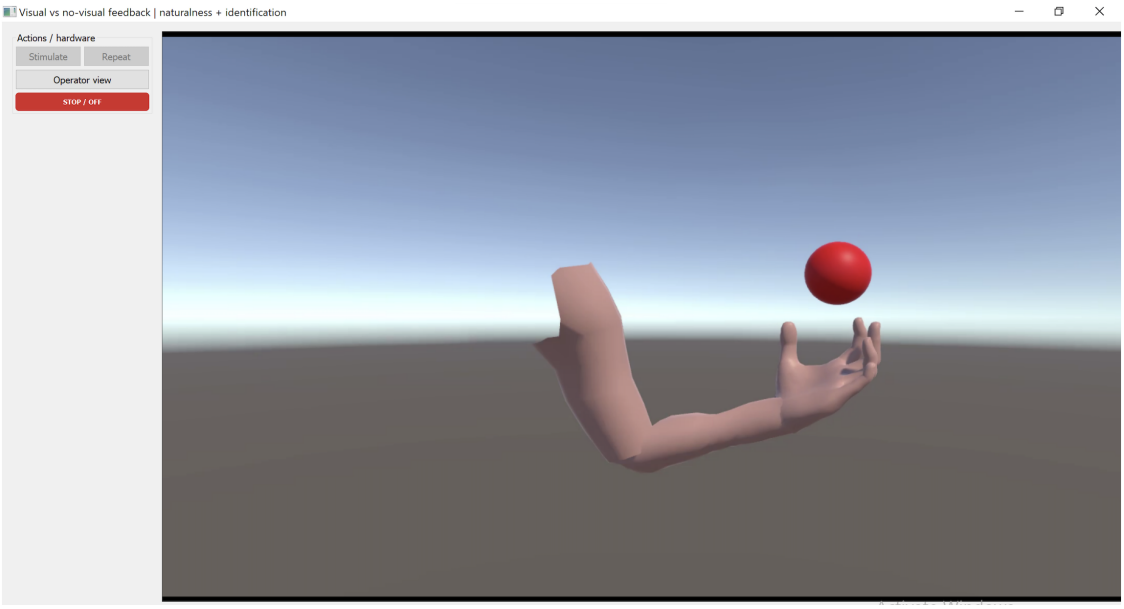


Figure 20: Stimulation window used during the visual feedback condition. The video of the simulated interaction is embedded in the window; in no visual feedback trials, the same display remains black.

Each participant completed 60 randomized trials, corresponding to:

$$6\,\text{situations} \times 5\,\text{encoding strategies} \times 2\,\text{feedback conditions} = 60\,\text{trials}.$$

The experiment was divided into two blocks of 30 trials. In visual feedback trials, stimulation was delivered together with the corresponding video of the simulated interaction. In no visual trials, stimulation was delivered without the video. After each trial, participants rated the perceived naturalness of the stimulation on a 1–5 scale. In no visual trials, they also identified which of the six situations the stimulation represented.

The six situations were intentionally selected to be partially distinguishable but not trivial. Since the stimulation interface provided only three spatial channels and projected sensory zones varied across subjects, the task was designed to produce intermediate identification performance and therefore allow comparison between encoding strategies.

4.6.3 Simulated tactile scenarios

The six simulated situations were:

- **Ball grasp:** the hand grasps a ball placed on the palm. Contact starts from the palm and becomes stronger as the fingers close around the object.
- **Ball slip:** the ball rolls across the palm, mainly from the radial to the ulnar side and back.
- **Ball bounce:** the interaction starts from a ball grasp like contact, but the ball is then launched into the air and caught again, producing repeated contact events.
- **Cone:** a single cone sequentially stimulates localized points of the palm.
- **Three cones:** three cone tips stimulate three localized palm regions simultaneously.
- **Finger tapping:** the fingers tap on a plane in a repeated sequence involving the thumb, the central fingers, and the little finger.

These scenarios were chosen to partially overlap in their spatial and temporal structure. For example, ball grasp and three cones can produce similar channel level pressure profiles, but differ in spatial extent: ball grasp is broader, whereas three cones are more localized. Similarly, cone stimulation and finger tapping both contain sequential localized events, while finger tapping and ball slip may both generate directional patterns across channels.

This makes the task informative: a strategy that only preserves the global pressure envelope may fail to distinguish some scenarios, whereas a more effective encoding should preserve enough spatial and temporal structure to improve identification. Because a real prosthetic hand equipped with pressure sensors was not available

in this work, the tactile inputs used by the encoding strategies were simulated. The hand was divided into 23 anatomical zones, and each zone was treated as if it were associated with one pressure sensor. The resulting synthetic pressure profiles therefore represented the temporal pressure that would have been measured by a distributed tactile sensor array during each simulated interaction.

Figure 21 shows the anatomical subdivision used to define these zone level pressure signals. The assignment of zones to the three stimulation channels was guided by the projected field mapping results reported in Section 5.1: thumb and radial palm zones were assigned to the radial channel, central finger and central palm zones to the medial channel, and little finger plus ulnar palm zones to the ulnar channel. This made the later encoding problem a compression problem, because rich spatial pressure information from 23 simulated zones had to be translated into a maximum of six stimulation commands: amplitude and frequency for each of the three available tTENS channels.

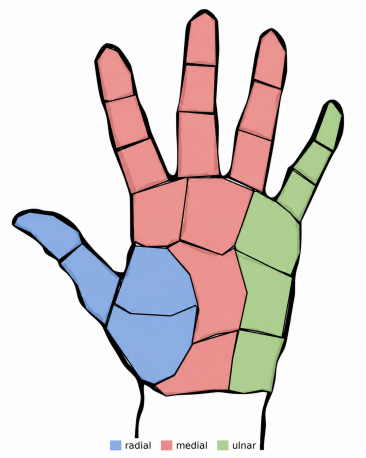


Figure 21: Schematic representation of the 23 anatomical hand zones used to simulate local pressure profiles.

4.6.4 Tested encoding strategies

Five encoding strategies were compared:

1. **Linear amplitude:** pressure was mapped directly to stimulation amplitude, while frequency was kept fixed at 50 Hz.
2. **Linear frequency:** pressure was mapped to stimulation frequency, while amplitude was fixed at 50% of the subject specific comfort range.
3. **TouchMime recruitment:** the TouchMime recruitment output controlled stimulation amplitude, while frequency was fixed at 50 Hz.

4. **TouchMime firing rate:** the TouchMime firing rate output controlled stimulation frequency, while amplitude followed the normalized pressure profile.

5. **SA/RA LIF:** simplified leaky integrate and fire neurons were driven by low pass and high pass filtered pressure signals, approximating slowly adapting and rapidly adapting tactile responses.

4.6.5 Linear amplitude and linear frequency strategies

The complete processing paths of the two linear strategies are summarized in Figure 22 using the ball-bounce scenario as an example.

The ball bounce pressure profiles used as a representative example are shown in Figure 23. These profiles are the same fixed tactile input used by all encoding strategies before strategy specific transformations are applied.

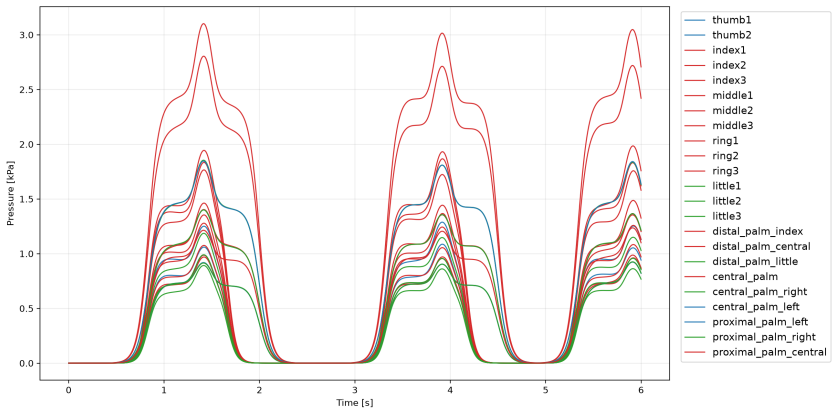


Figure 23: Example of 23 zone pressure profiles for the ball bounce scenario. Each curve represents the synthetic pressure evolution of one anatomical hand zone.

For the two linear baseline strategies, however, the stimulation interface has only three output channels:

- **CH1:** radial, thumb and thenar side.
- **CH2:** medial, index, middle, ring and central palm.
- **CH3:** ulnar, little finger and hypothenar side.

This creates an information compression problem: the 23 simulated pressure profiles contain richer spatial information than the stimulation hardware can directly reproduce, while the final command space contains at most six time varying parameters, namely amplitude and frequency for each of the three channels.

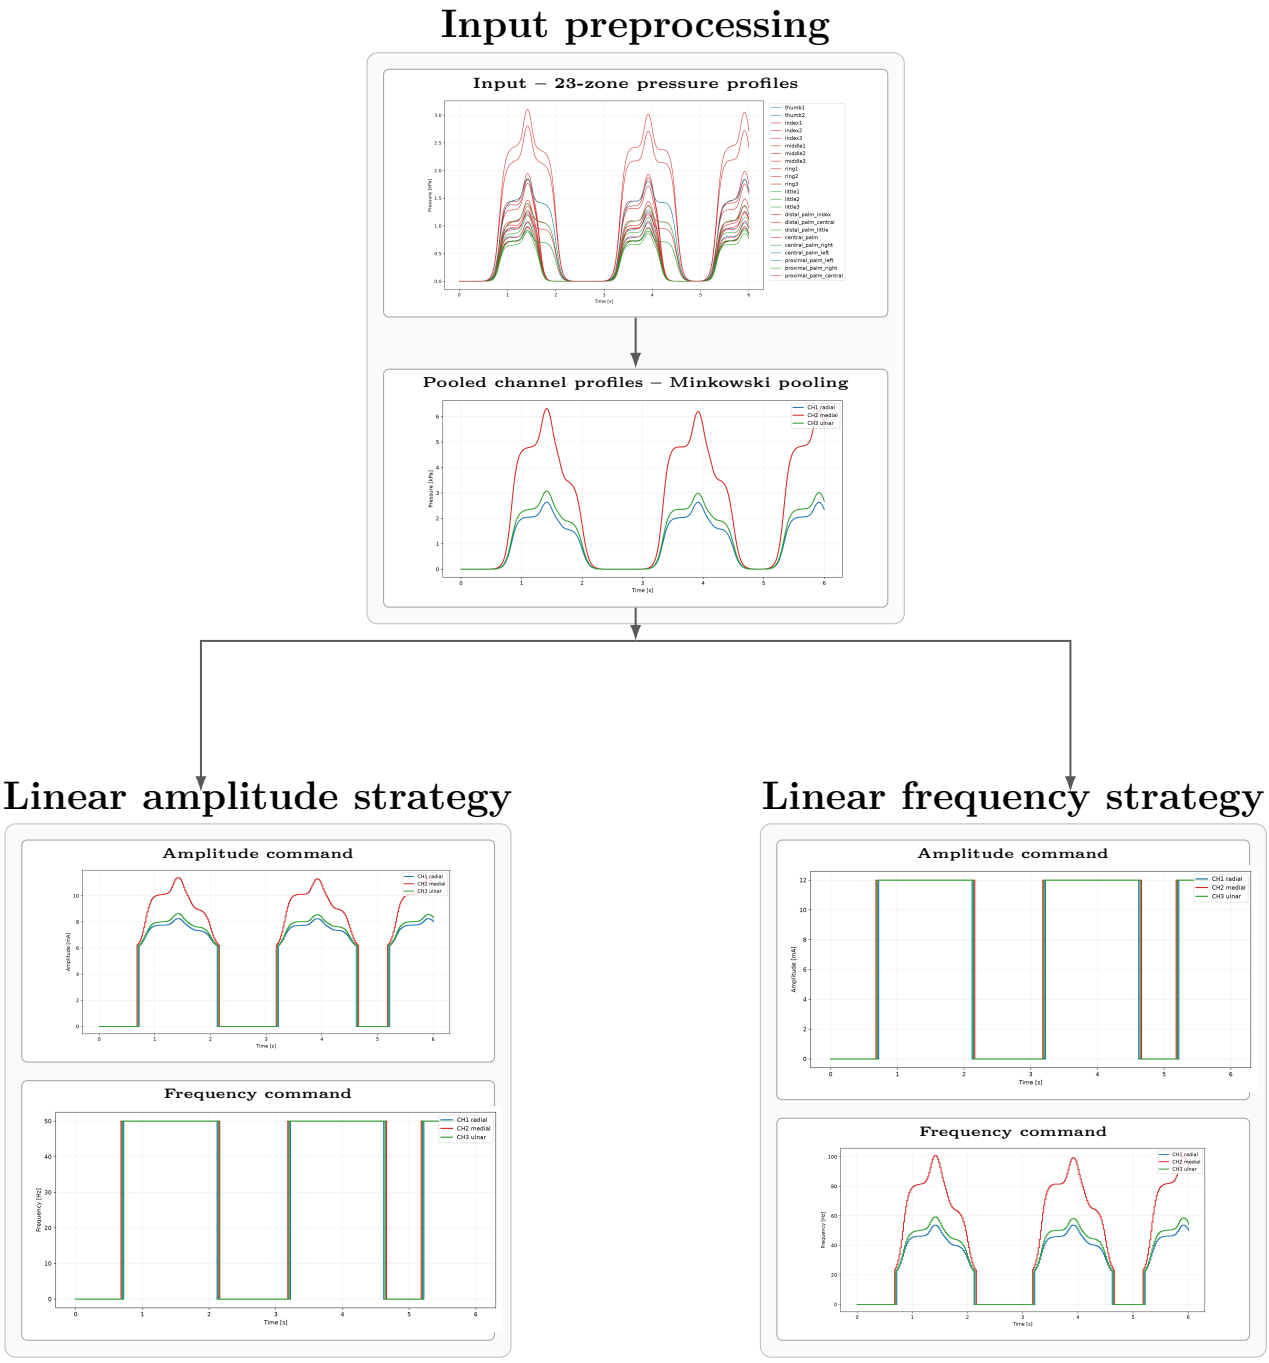


Figure 22: Overview of the linear encoding pipelines for the ball-bounce scenario. The 23-zone pressure input is first pooled into channel level profiles using Minkowski pooling. The pooled channel signals are then mapped either through the linear amplitude strategy or through the linear frequency strategy, producing the corresponding stimulation amplitude and frequency commands.

Moreover, the number of zones assigned to each channel is not the same; the medial channel contains more anatomical regions than the radial and ulnar channels.

A simple sum of local pressures would therefore make channels with more zones artificially larger, even if the local contact intensity were not stronger. Conversely, averaging over all zones in a channel would penalize localized events, such as cone contact, especially in broad channels.

To address this, zone level pressures were aggregated using a Minkowski like compressed spatial pooling rule:

$$D_{ch}(t) = \left(\sum_{z \in ch} p_z(t)^q \right)^{1/q}, \qquad q = 2,$$

where $p_z(t)$ is the synthetic pressure profile of zone z, and $D_{ch}(t)$ is the resulting pressure drive for stimulation channel ch. In the current implementation, $q = 2$:

$$D_{ch}(t) = \sqrt{\sum_{z \in ch} p_z(t)^2}.$$

This rule behaves as a compromise between maximum pooling and linear summation. If multiple neighbouring zones are active, the channel drive increases, but less strongly than with a direct sum. For example, if three zones have the same local pressure, the pooled drive increases by a factor of $\sqrt{3}$ rather than by a factor of 3.

This was useful for the present application because it preserves the contribution of simultaneously active zones while preventing broad channels from dominating only because they contain more simulated regions. Minkowski summation is used in haptic psychophysics to describe how multiple spatial or sensory contributions combine into an overall response: the exponent controls the transition between averaging-like integration and linear addition. This makes it a relevant engineering choice for the present spatial-compression problem. It is not, however, treated here as a validated physiological model of tactile perception.

Figure 24 shows the resulting three channel pressure drive after Minkowski pooling for the ball bounce example.

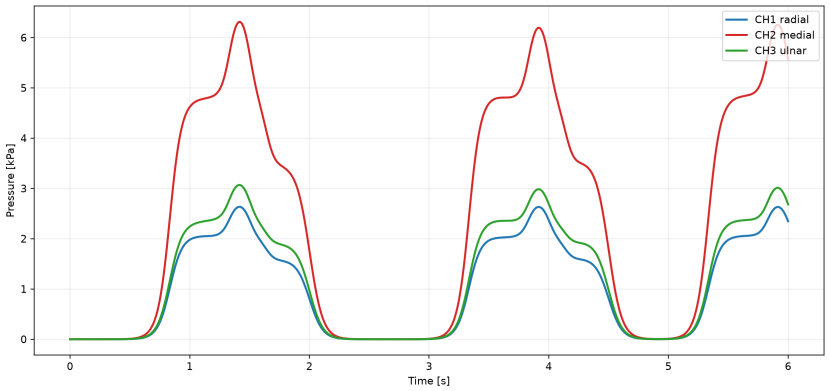


Figure 24: Three channel pressure drive obtained after Minkowski pooling of the 23 zone pressure profiles for the ball bounce scenario.

The resulting three channel pooled pressure drive was used only for the two linear baseline strategies. In the linear amplitude strategy, it directly controlled stimulation amplitude; in the linear frequency strategy, it directly controlled stimulation frequency. In the TouchMime and SA/RA LIF strategies, instead, the model operated on zone level or region level pressure profiles before channel level aggregation, as described in the following sections.

For the two linear baseline strategies, the pooled channel drive was first normalized before being mapped onto stimulation parameters:

$$\widehat{D}_{ch}(t) = \min\left(\frac{D_{ch}(t)}{P_{ref}}, 1\right).$$

Here, P_{ref} is the pressure reference used to define full scale activation. In the implemented GUI, this value was computed as the global maximum of the available Minkowski pooled channel pressure profiles across the simulated scenario set. Therefore, full scale activation means that the strongest pooled pressure drive observed in the reference scenario set is mapped to a normalized value of 1, while lower pressure values are expressed as fractions of this reference.

In the linear amplitude strategy, the normalized pressure drive controlled stimulation amplitude within the subject specific comfort range, while the stimulation frequency was kept fixed at $f_{ref} = 50\,\text{Hz}$:

$$A_{ch}(t) = T_{\min,ch} + \min\left(\frac{D_{ch}(t)}{P_{ref}}, 1\right)(T_{\max,ch} - T_{\min,ch}).$$

In the linear frequency strategy, amplitude was fixed at 50% of the subject specific comfort range and the normalized pressure drive controlled stimulation frequency:

$$f_{ch}(t) = f_{\min} + \min\left(\frac{D_{ch}(t)}{P_{ref}}, 1\right)(f_{\max} - f_{\min}).$$

Therefore, P_{ref} defines how pressure is normalized before the final mapping into amplitude or frequency commands.

4.6.6 TouchMime encoding strategy

The complete TouchMime processing path is summarized in Figure 25 using the ball-bounce scenario as an example.

TouchMime is the most biomimetic encoding model used in this work. Rather than mapping pressure directly to stimulation variables, it was used as a computationally lighter surrogate of TouchSim. TouchSim predicts the spike trains generated by populations of SA1, RA, and PC tactile afferent fibres in response to an indentation applied at a specified location on the glabrous hand. From these simulated population responses, two aggregate descriptors can be extracted: the population firing rate, which describes the mean spiking activity of the activated afferents over time, and the active area, which describes the spatial extent of the afferent population recruited by the stimulus. TouchMime approximates these two outputs directly from the indentation profile through regression models, avoiding the explicit simulation of thousands of individual afferent spike trains. In the present work, these outputs were then used as candidate control variables for stimulation frequency and stimulation amplitude, respectively.

The TouchMime models were previously identified by simulating circular indentation stimuli in anatomically defined hand regions, sampled at 100 Hz and using a lag parameter equal to 3. The 100 Hz sampling rate represents one indentation sample every 10 ms. The lag parameter specifies how many recent samples of indentation position, speed, and acceleration are included in the regression input; a value of 3 therefore provides the model with a short 30 ms history of the stimulus, allowing it to capture temporal dynamics rather than relying only on the instantaneous indentation. Therefore, in the online encoding stage, the original 23 simulated pressure profiles were first reorganized into the 13 circular anatomical regions used by TouchMime. These regions were chosen to reproduce the contact layouts of the simulated scenarios. Figure 26 shows the TouchMime regions used for the ball related scenarios and for the localized cones condition.

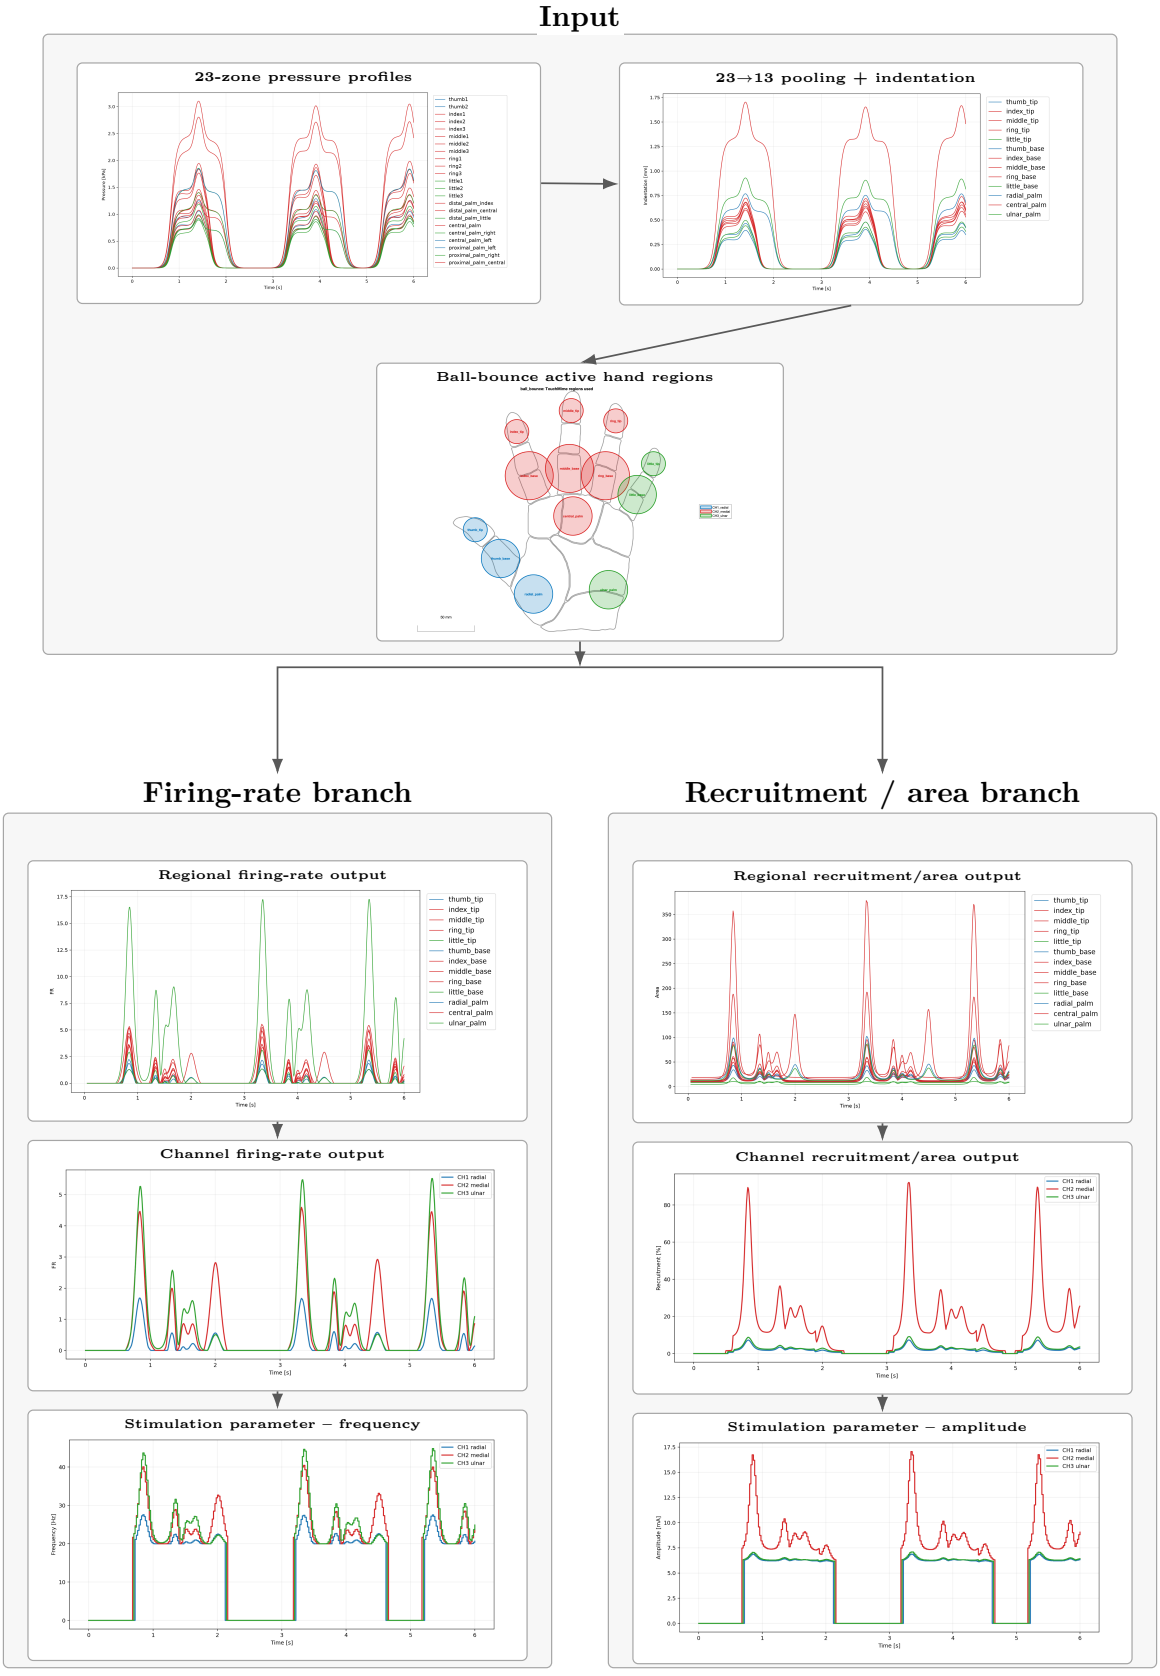


Figure 25: TouchMime-based encoding pipeline for the ball-bounce scenario. The 23-zone pressure profiles are pooled into 13 regions, converted into indentation inputs, processed by TouchMime, and mapped to frequency and amplitude stimulation commands.

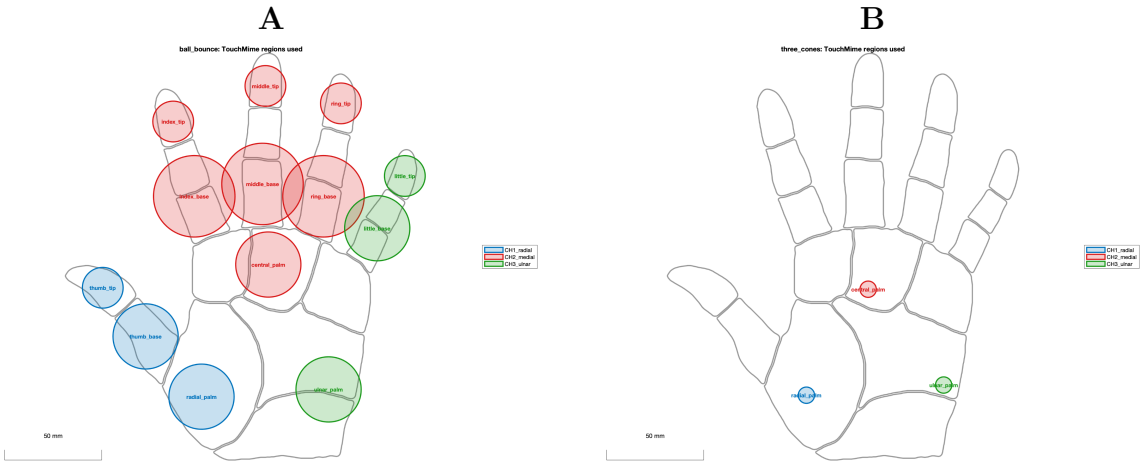


Figure 26: TouchMime anatomical regions used to model the simulated contact scenarios. (A) Broad set of regions used for the ball bounce example; the same spatial layout was also used for ball grasp and ball slip, with different temporal activation profiles. (B) Localized regions used for the three cones example.

This intermediate representation was preferred to a direct 23 region TouchMime implementation because the final stimulation interface had only three output channels and because the TouchMime Area output already represents a spatially extended recruited afferent territory. If many neighbouring local zones were treated as fully independent and their Area outputs were summed afterwards, channel recruitment could be overestimated. The adopted compromise was therefore:

$$\begin{aligned} &23 \text{ pressure zones} \rightarrow 13 \text{ anatomical regions} \rightarrow \text{TouchMime} \\ &\qquad\qquad \rightarrow 3 \text{ TENS stimulation channels.} \end{aligned}$$

Before the TouchMime regression step, the 23 local pressure profiles were pooled into the 13 anatomical regions by Minkowski pooling with exponent $q = 2$:

$$p_r(t) = \left(\sum_{z \in r} p_z(t)^q \right)^{1/q}, \qquad q = 2,$$

where $p_r(t)$ is the pooled pressure profile of region r. This pooling preserves the effect of multiple simultaneously active neighbouring zones without resorting to a purely linear sum. A single global pressure reference P_{99} was then computed as the 99th percentile of all positive samples from these pooled regional pressure profiles across all simulated scenarios. Therefore, the normalization was not performed separately for each region or for each scenario: the same global reference was used for all TouchMime inputs so that relative differences in intensity across gestures were preserved. The conversion from pressure to indentation was essential because TouchMime accepts indentation profiles, rather than pressure profiles, as its input. Each pooled pressure profile was then remapped linearly to an indentation like

input according to

$$d_r(t) = \min\left(d_{\max},\, d_{\max}\frac{p_r(t)}{P_{99}}\right),$$

with $d_{\max} = 3$ mm. This value should be interpreted as an engineering scaling limit rather than as a calibrated biomechanical parameter: it bounds the synthetic pressure to indentation conversion within a plausible tactile range and prevents large pressure peaks from dominating the TouchMime input. Figure 27 shows the resulting regional indentation profiles for the ball bounce example.

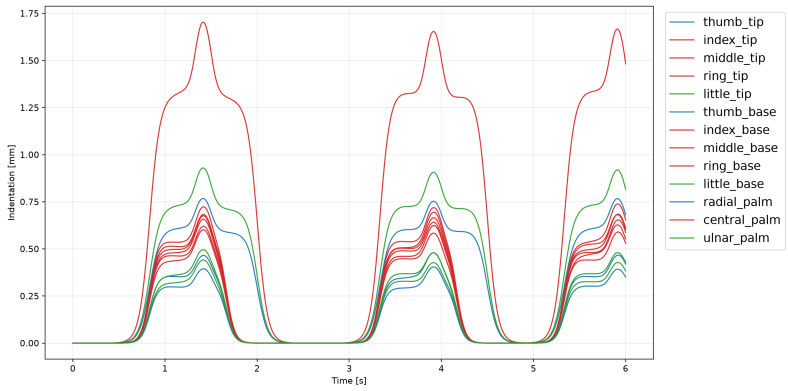


Figure 27: Regional indentation profiles used as TouchMime input for the ball bounce example.

For each regional indentation input, TouchMime predicted two outputs: a firing rate signal $FR_r(t)$ and an Area signal $Area_r(t)$.

In the present work, $FR_r(t)$ was used as a dynamic descriptor for frequency modulation, whereas $Area_r(t)$ was used as a recruitment descriptor for amplitude modulation. Figures 28 and 29 show these regional outputs for the ball bounce example.

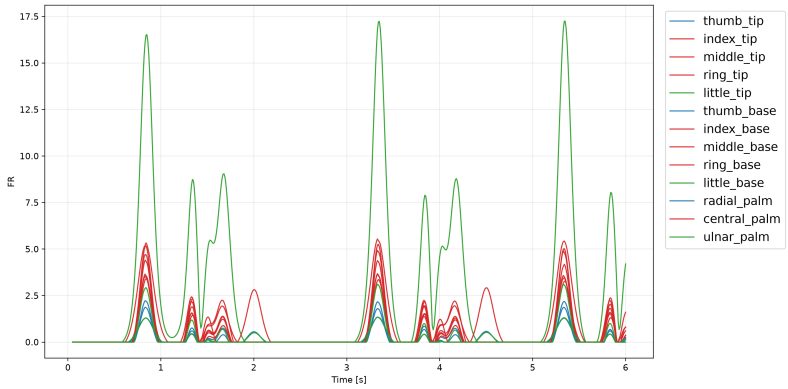


Figure 28: Regional TouchMime firing rate outputs for the ball bounce example.

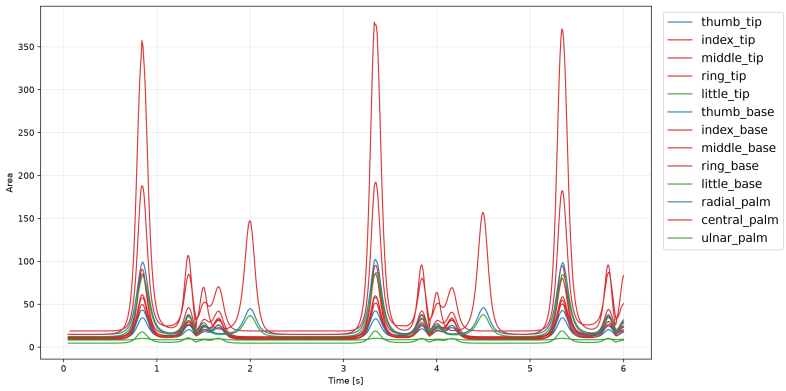


Figure 29: Regional TouchMime Area outputs for the ball bounce example.

To obtain the final stimulation commands, the TouchMime outputs were aggregated into the three stimulation channels. In the current implementation, channel firing rate was computed as a simple mean over the active regions belonging to the same channel:

$$FR_{ch}(t) = \frac{1}{N_{\text{active},ch}(t)} \sum_{r \in ch,\, active} FR_r(t).$$

This choice was adopted because channel firing rate was intended to represent the average dynamic response of the active contact regions, without increasing automatically just because more regions were active. In contrast, the Area output was treated as an extensive spatial quantity. The Area signals of the regions belonging to the same channel were summed and normalized by an empirical channel capacity, defined as the sum of the maximum positive Area values observed for those regions across all simulated scenarios. The resulting channel recruitment was therefore

$$Recruitment_{ch}(t) = 100\, \frac{\sum_{r \in ch} Area_r(t)}{\sum_{r \in ch} Capacity_r}.$$

Figure 30 shows the channel level TouchMime firing rate output, while Figure 31 shows the corresponding channel recruitment obtained from the normalized Area signals.

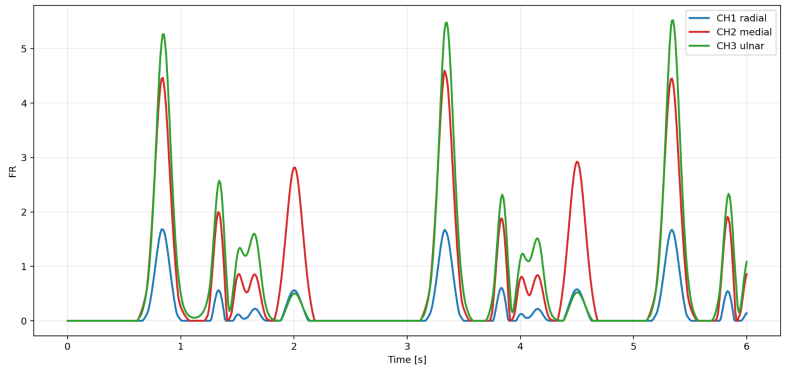


Figure 30: Channel level TouchMime firing rate output for the ball bounce example.

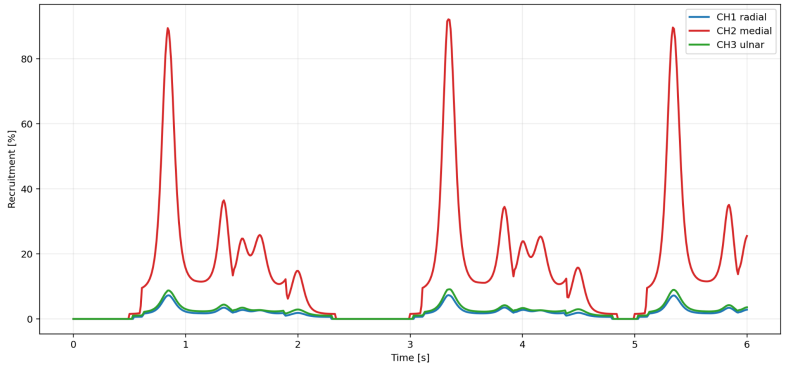


Figure 31: Channel level recruitment obtained from the TouchMime Area output for the ball bounce example.

Because recruitment sums the Area contributions of simultaneously active regions, the medial channel can dominate in scenarios such as ball bounce, where several medial regions are recruited. This was intentional: recruitment encoded the spatial extent of activation, consistently with the mapping experiment, in which increasing stimulation amplitude generally expanded the projected sensory area.

Two TouchMime based stimulation strategies were then tested. In the TouchMime firing rate strategy, the channel firing rate controlled stimulation frequency through the linear mapping

$$f_{ch}(t) = f_{\min} + \min\left(\frac{FR_{ch}(t)}{FR_{ref}}, 1\right)(f_{\max} - f_{\min}),$$

where FR_{ref} is the firing rate reference corresponding to the maximum stimulation frequency. The resulting frequency command for the ball bounce example is shown in Figure 32. In the TouchMime recruitment strategy, the normalized recruitment controlled stimulation amplitude within the subject specific comfort range, while stimulation frequency was kept fixed at 50 Hz. The corresponding

amplitude command was computed as

$$A_{ch}(t) = T_{\min,ch} + \frac{Recruitment_{ch}(t)}{100} \left(T_{\max,ch} - T_{\min,ch}\right),$$

where $T_{\min,ch}$ and $T_{\max,ch}$ are the detection and maximum comfortable thresholds of channel ch. For the illustrative example reported here, a common comfort range of 6–18 mA was used for all channels. The resulting amplitude command is shown in Figure 33.

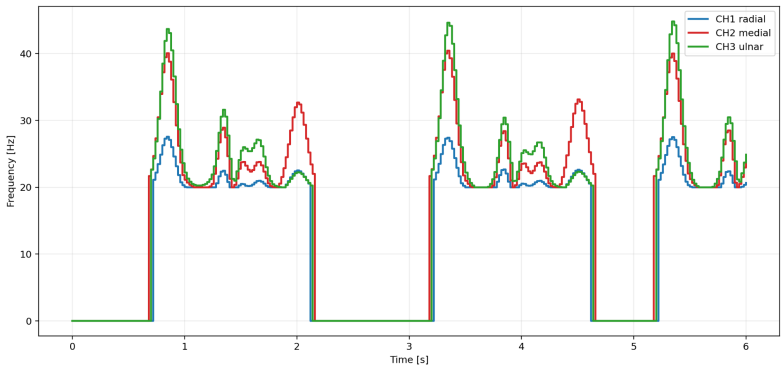


Figure 32: Channel level stimulation frequency obtained from the TouchMime firing rate strategy for the ball bounce example.

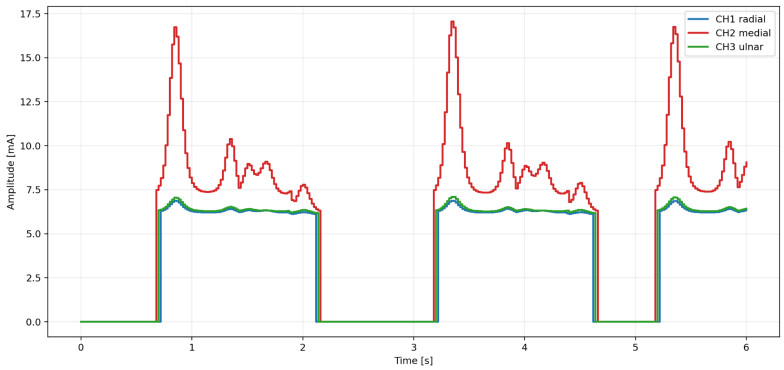


Figure 33: Channel level stimulation amplitude obtained from the TouchMime recruitment strategy for the ball bounce example.

Overall, TouchMime was the most biomimetic strategy tested, but remained an engineering approximation. Its synthetic pressure inputs, globally scaled pressure to indentation conversion, and compression from 23 zones to three stimulation channels were practical design choices rather than validated models of tactile perception or transcutaneous stimulation.

4.6.7 SA/RA LIF encoding strategy

The complete SA/RA LIF processing path is summarized in Figure 34 using the ball-bounce scenario as an example. The SA/RA LIF strategy was implemented as a biomimetic encoding approach designed to preserve a key biological principle, namely the separation between sustained pressure information and transient contact events, while remaining simpler than a full tactile afferent model such as TouchMime. Its main objective is to transform synthetic hand pressure profiles into a firing rate signal and to remap this output to stimulation frequency. Stimulation amplitude is not generated by the LIF model itself; instead, it is mapped linearly from the pooled channel pressure using the same pressure to comfort range rule adopted for the linear amplitude baseline.

The complete processing chain is illustrated using the ball bounce scenario as a representative example. It starts from the same 23 zone pressure profiles introduced in Figure 23, which represent the fixed tactile input used by all strategies before strategy specific processing.

For each of the 23 anatomical hand zones, the local pressure signal is first normalized with respect to a global pressure reference. This avoids making the LIF dynamics dependent on the absolute scale of the synthetic pressure profiles. The normalized signal is then split into two parallel branches to approximate the behaviour of slowly adapting and rapidly adapting tactile afferents. The SA like branch, shown in Figure 35, is obtained through second order Butterworth low pass filtering with a cutoff frequency of 0.7 Hz and captures the sustained component of contact. The RA like branch, shown in Figure 36, is obtained through second order Butterworth high pass filtering with the same cutoff frequency of 0.7 Hz, followed by rectification, so that both increases and decreases in pressure generate positive transient activity.

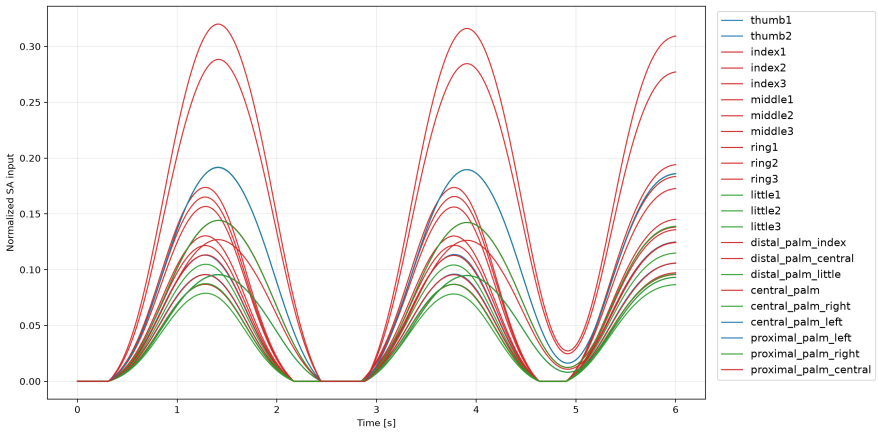


Figure 35: SA like low pass filtered zone-level input for the ball-bounce example.

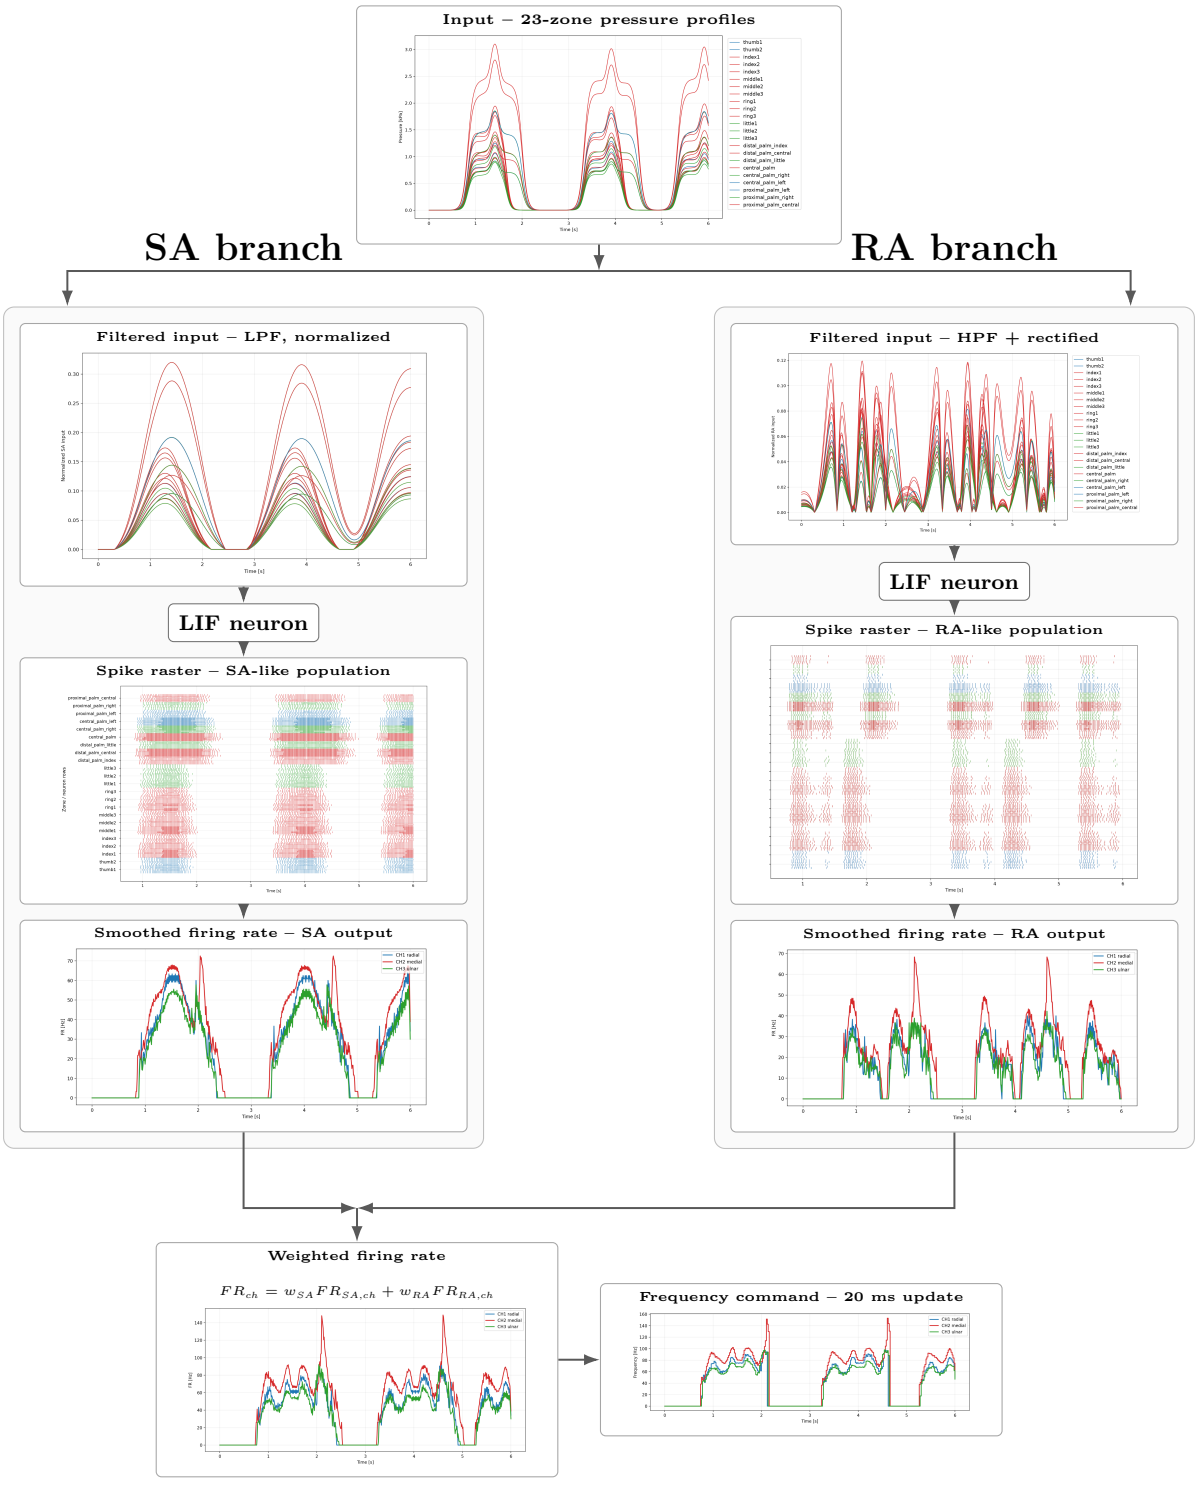


Figure 34: Overview of the SA/RA LIF encoding pipeline for the ball-bounce scenario. The 23-zone pressure input is processed through a slowly adapting branch and a rapidly adapting branch, converted into spike activity, smoothed into channel level firing rates, and finally mapped to the stimulation-frequency command.

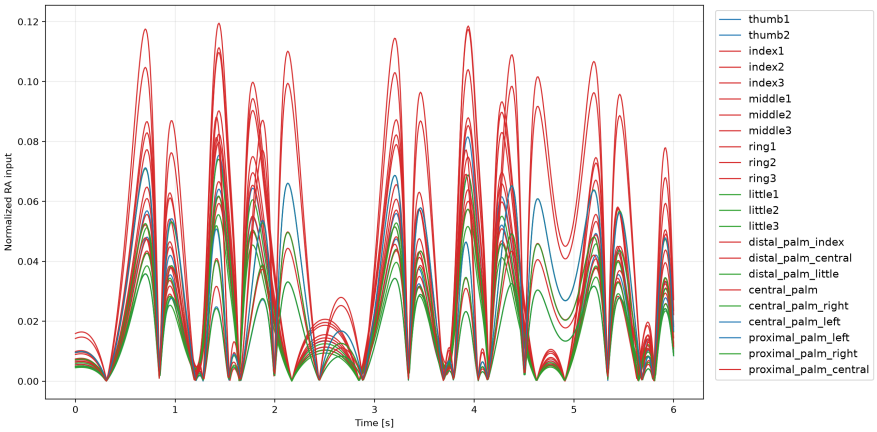


Figure 36: RA like high pass filtered and rectified zone-level input for the ball-bounce example.

Each filtered signal then drives a leaky integrate and fire (LIF) unit, a simplified neuron model that represents how a biological neuron accumulates incoming current in its membrane potential while part of that potential gradually leaks away. When the accumulated potential reaches a threshold, the model emits a discrete spike, resets, and remains temporarily inactive, reproducing the basic integrate–fire–refractory sequence without modelling detailed ion-channel dynamics. In this implementation, leakage is controlled by the membrane time constant, set to $\tau_m = 40\,\text{ms}$. When the membrane potential crosses the spike threshold, $V_{th} = 0.20$, a spike is generated, the membrane is reset, and the unit remains inactive during a refractory period of 5 ms. The resulting spike trains are shown for the SA and RA branches in Figures 37 and 38.

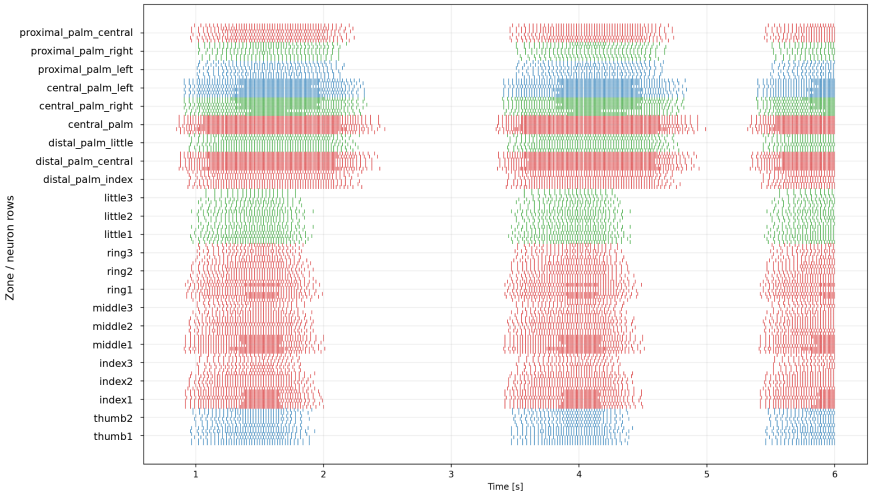


Figure 37: Spike raster generated by the SA like LIF population.

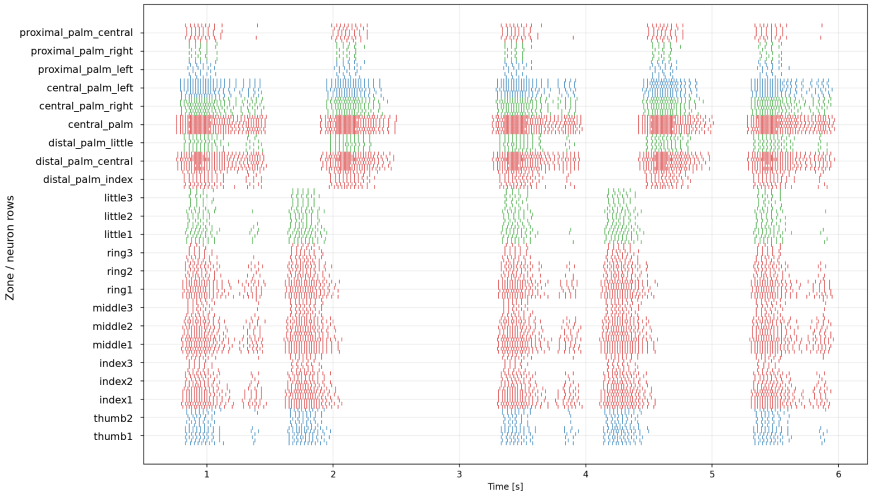


Figure 38: Spike raster generated by the RA like LIF population.

The spike trains are then converted into firing rate profiles. Since the raw spike derived rate can fluctuate strongly at the sampling resolution of the pressure profiles, the firing rate is smoothed using a moving average window of 50 ms. The pressure profiles are sampled at 100 Hz, corresponding to a 10 ms temporal resolution. The smoothed channel level outputs of the SA and RA branches are shown in Figures 39 and 40, respectively.

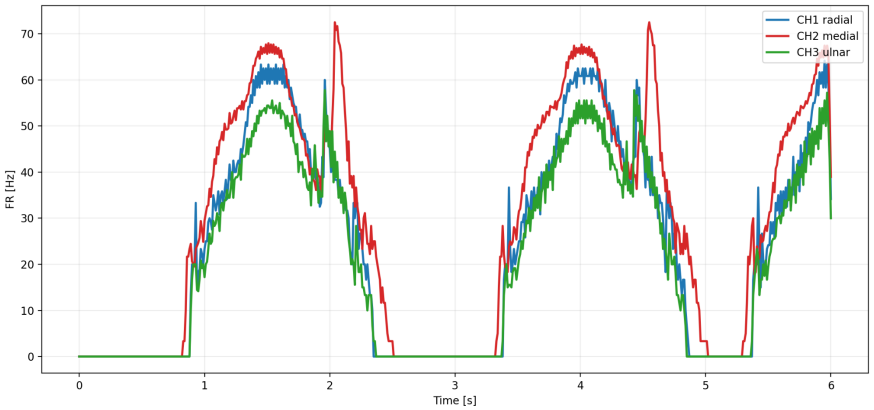


Figure 39: Smoothed channel-level firing rate obtained from the SA branch.

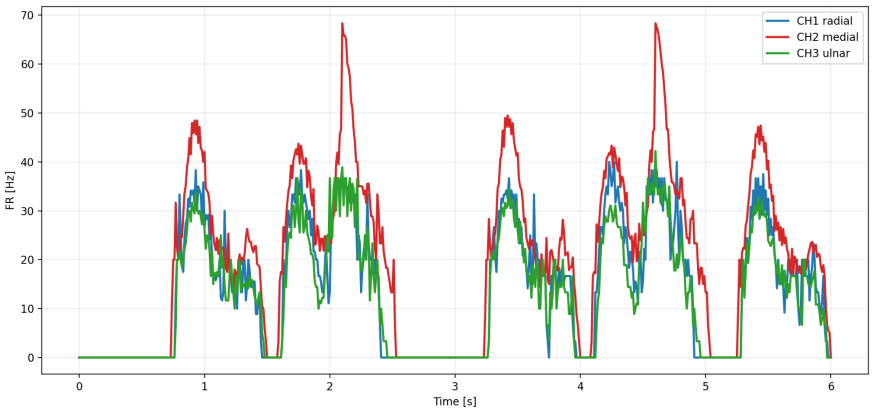


Figure 40: Smoothed channel-level firing rate obtained from the RA branch.

The smoothed SA and RA firing rates are then combined through output weights. The SA weight was set to $w_{SA} = 1.0$ to preserve the sustained contact component, while the RA weight was set to $w_{RA} = 1.3$ to slightly emphasize transient events. After this combination, zone level activity is aggregated into the three stimulation channels by averaging only the active zones within each channel. This avoids diluting the channel response with inactive zones, particularly for the medial channel, which contains more anatomical regions than the radial and ulnar channels.

The weighted channel firing rate, shown in Figure 41, is computed as:

$$FR_{ch}(t) = w_{SA} FR_{SA,ch}(t) + w_{RA} FR_{RA,ch}(t).$$

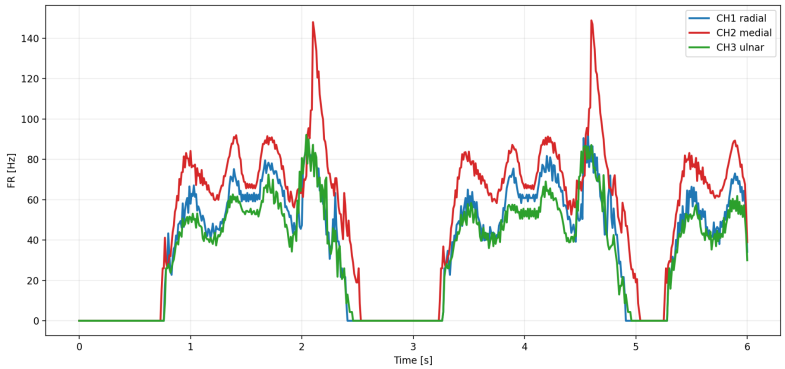


Figure 41: Weighted SA/RA channel firing rate after combining the sustained and transient components.

The final channel firing rate is mapped linearly onto stimulation frequency. The minimum and maximum stimulation frequencies were set to $f_{\min} = 20\,\text{Hz}$ and $f_{\max} = 200\,\text{Hz}$, respectively. A rate reference parameter, FR_{ref}, defines the firing rate that corresponds to the maximum stimulation frequency; increasing this

reference reduces the delivered frequency for the same model output and therefore acts as a final scaling parameter. The mapping is:

$$f_{ch}(t) = \begin{cases} f_{\min} + \min\left(\dfrac{FR_{ch}(t)}{FR_{ref}}, 1\right)(f_{\max} - f_{\min}), & FR_{ch}(t) > 0, \\ 0, & FR_{ch}(t) = 0. \end{cases}$$

The firing rate to frequency mapping was first computed on the original pressure profile time base and then resampled to the 20 ms command update window used by the Arduino scheduler. Before being sent to the hardware, the command was gated by contact related pressure information: when the corresponding channel pressure was below the contact threshold, both amplitude and frequency were set to zero, avoiding stimulation during no contact periods. In this sense, the final frequency command is the LIF derived frequency profile after resampling and contact gating, as shown in Figure 42.

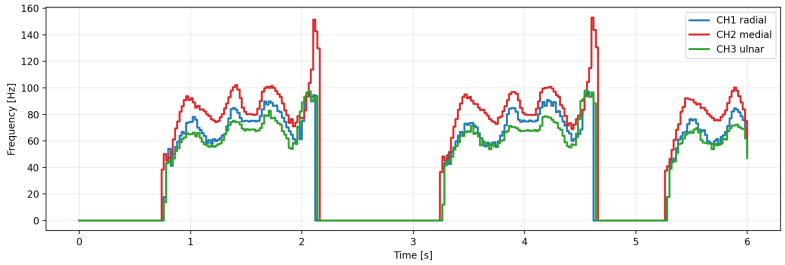


Figure 42: Final frequency command obtained from the SA/RA LIF firing rate after 20 ms resampling and contact gating. This is the frequency profile sent to the stimulation hardware.

The filter cutoffs, LIF parameters, smoothing window, command update interval, branch weights, and frequency limits described above were selected after preliminary tuning on a few pilot subjects and were kept fixed during the experimental comparison. The firing-rate reference FR_{ref} was the only final scaling term adjusted through the load-matching procedure described below.

4.6.8 Subject specific scaling and load matching

All amplitude commands were scaled within subject specific comfort ranges. For each stimulation channel, the detection threshold $T_{\min,ch}$ and the maximum comfortable threshold $T_{\max,ch}$ were obtained from previous calibration experiments, defining the available amplitude range for that subject and channel. This was necessary because the same stimulation current can be perceived very differently across subjects, channels, and electrode positions.

In addition, load matching was applied across encoding strategies. Here, load matching means adjusting one overall scale parameter for each strategy so that

its time-integrated proxy for delivered electrical charge is comparable with that of the reference strategy, while preserving the temporal shape and relative channel pattern produced by the encoder. This was introduced to reduce differences in the total amount of charge delivered to the subject. Without this step, a strategy could appear more or less natural simply because it delivered a higher stimulation load, rather than because its temporal or biomimetic structure was more effective. The load matching procedure should therefore be interpreted as a practical control step for comparing encoding strategies more fairly.

The matching criterion was based on a charge rate load proxy. For each time sample, the instantaneous charge rate is proportional to stimulation amplitude, pulse width and stimulation frequency. The integrated load used for matching was therefore computed as

$$L = \sum_{ch=1}^{3} \sum_{k} A_{ch}(t_k)\, PW\, f_{ch}(t_k)\, \Delta t,$$

where k indexes the discrete update samples, $A_{ch}(t_k)$ is the absolute stimulation amplitude, PW is the pulse width, $f_{ch}(t_k)$ is the stimulation frequency, and Δt is the update period. Because PW was fixed across all strategies, it does not affect the relative matching, but it is included in the formula to make explicit that L is proportional to the charge rate delivered to the subject over the stimulation duration.

The linear amplitude strategy was used as the reference condition and was left unchanged. For each other strategy, a single scaling parameter was adjusted while preserving the temporal shape of the encoded profile:

- **Linear frequency:** the pressure reference p_{ref} was adjusted in the pressure to frequency mapping.
- **TouchMime recruitment amplitude:** the recruitment gain g_{rec} was adjusted in the amplitude mapping.

  $$A_{ch}(t) = T_{\min,ch} + \mathrm{clip}\left(g_{rec}\frac{Recruitment_{ch}(t)}{100}, 0, 1\right)\left(T_{\max,ch} - T_{\min,ch}\right).$$

  This parameter was chosen because the TouchMime recruitment output is a normalized model derived descriptor rather than an absolute physiological recruitment fraction.
- **TouchMime firing rate frequency:** the firing rate reference FR_{ref} was adjusted in the final frequency mapping.
- **SA/RA LIF:** the channel firing rate reference FR_{ref} was adjusted in the final firing rate to stimulation frequency mapping. This changed the output frequency scale without retuning the internal SA/RA or LIF dynamics.

5 Results and discussion

The results are organized according to the experimental pipeline. This separates the description of each protocol from the corresponding outcome figures and interpretation.

5.1 Threshold estimation and projected sensation mapping

The primary purpose of the first experiment was to calibrate the usable stimulation space for each subject and channel. The detection threshold and maximum comfortable threshold defined the subject specific comfort range within which stimulation amplitude could subsequently be modulated. Figure 43 summarizes these ranges for all 10 participants. Both limits varied substantially across subjects and channels, showing that a common fixed operating range would not have provided comparable or consistently comfortable stimulation.

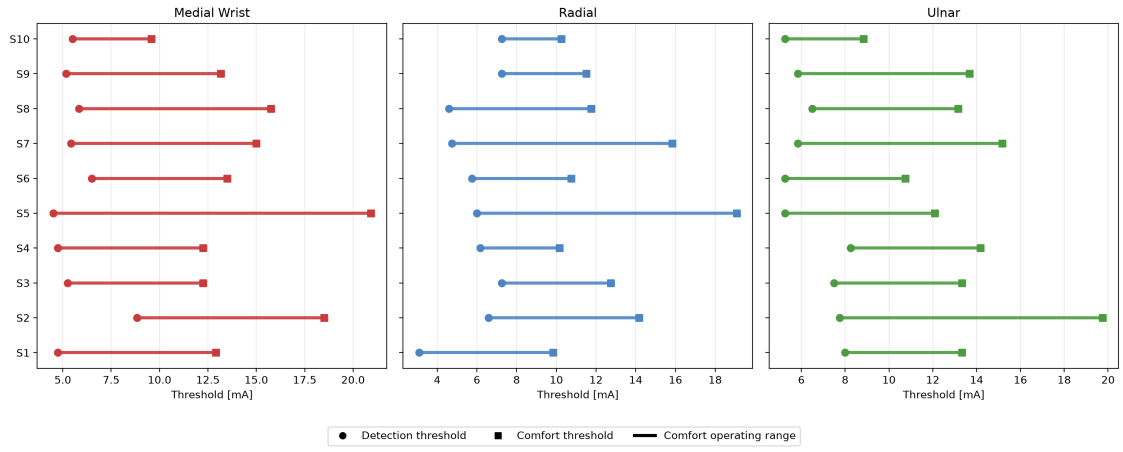


Figure 43: Comfort operating ranges across subjects for the medial wrist, radial, and ulnar channels.

During the maximum threshold staircase, the projected hand zones were recorded at every 0.5 mA amplitude step. For clarity, Figure 44 shows four representative levels for one subject: $T_{\min}$, 33% and 66% of the comfort range, and $T_{\max}$. Increasing current frequently enlarged the projected field in addition to increasing perceived intensity. At the lowest level, a clearly coloured referred zone was not always reproduced even though $T_{\min}$ had been defined as the onset of referred sensation. This is consistent with short term adaptation, attention, and small changes in the electrode–skin interface, and indicates that the threshold is an operational estimate rather than a perfectly sharp perceptual boundary.

The medial channel shows the clearest expansion in the representative subject and generally recruited the broadest territory at group level, often extending over the central palm and fingers. Partial overlap between radial and medial projected

fields was expected: both electrode placements targeted different regions of the median nerve territory, so their main fields differed but were not anatomically independent.

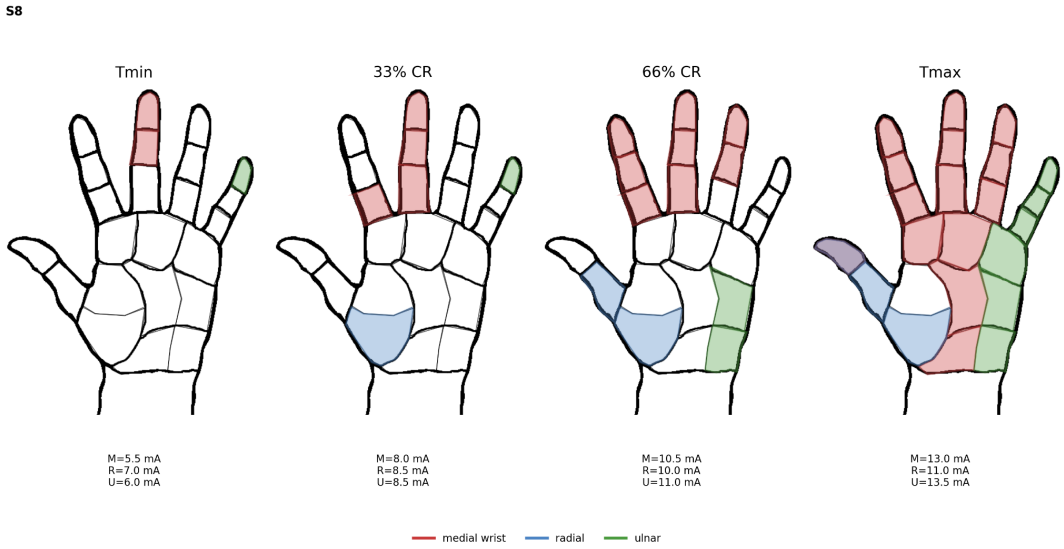


Figure 44: Projected field evolution with amplitude in one representative subject.

Figure 45 shows the corresponding group level evolution. Darker zones indicate regions reported by more participants. Despite inter-subject variability, the field generally expanded with amplitude, particularly for the medial channel.

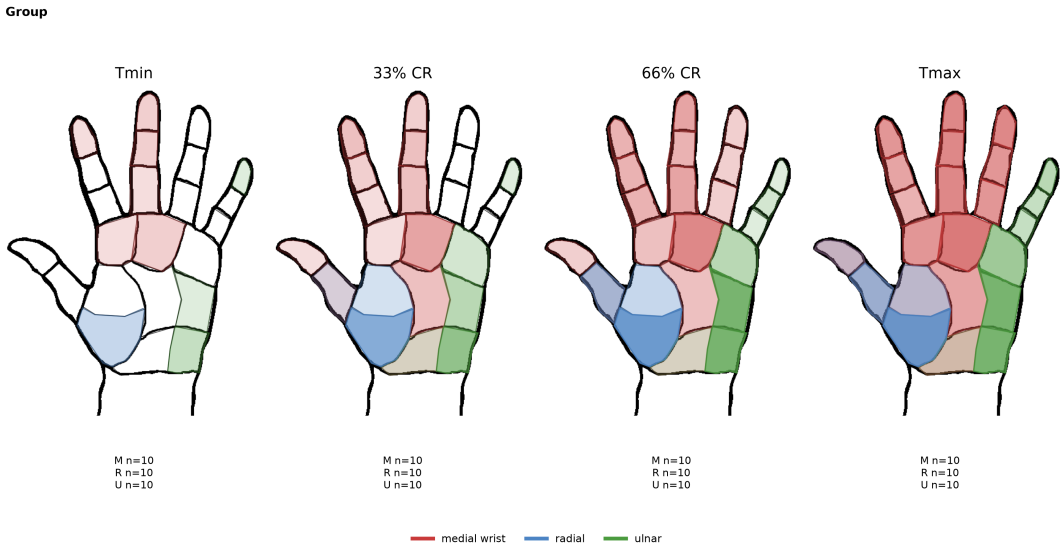


Figure 45: Group level projected field evolution with amplitude.

A group summary of the projected fields is shown alongside the functional three-channel hand division in Figure 46. The medial channel produced the broadest

field, the radial channel mainly covered the thumb and radial palm, and the ulnar channel was concentrated on the little finger and ulnar palm. These experimentally observed territories guided the later assignment of the 23 simulated pressure zones to the three stimulation channels. The assignment should therefore be interpreted as a functional compression map rather than as an exact anatomical separation of the nerves.

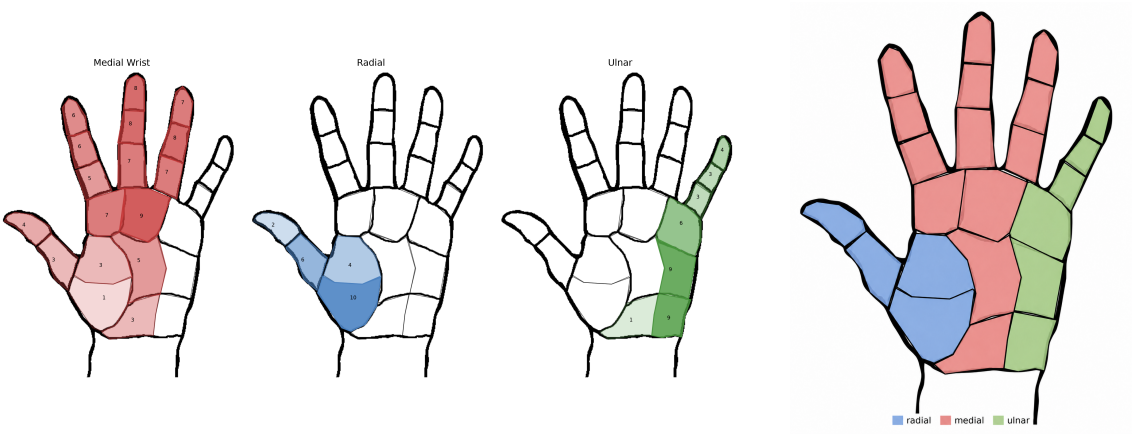


Figure 46: Group level projected fields and the resulting functional division of the hand zones across stimulation channels.

The referred sensation balance was more variable than the spatial extent of the projected field. Figure 47 shows three different behaviours in the same representative subject: the radial score increases, the medial score remains high before decreasing, and the ulnar score remains mostly stable. The group profiles in Figure 48 are comparatively flat and have wide variability bands, confirming that no common monotonic relationship with amplitude can be identified across subjects. The ulnar channel is the only one showing a weak recurring tendency to start from lower referred scores and then increase or remain stable. This is consistent with the practical observation that, particularly at low amplitudes, the ulnar site was often the most difficult channel in which to evoke a clear hand sensation rather than a sensation mainly under the electrode.

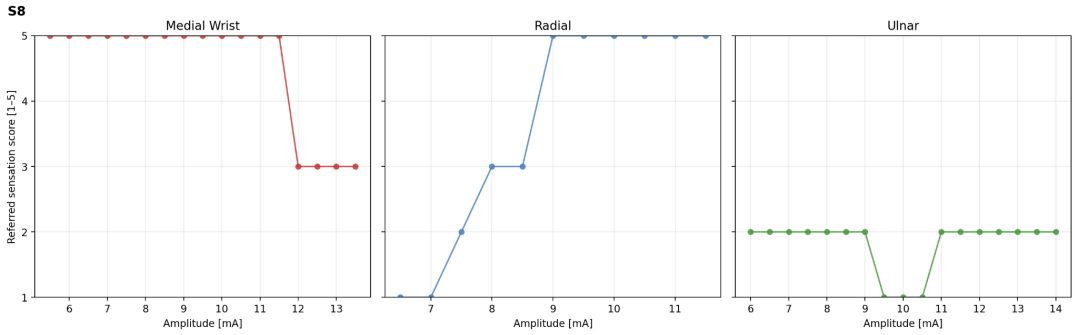


Figure 47: Referred sensation balance versus amplitude in one representative subject.

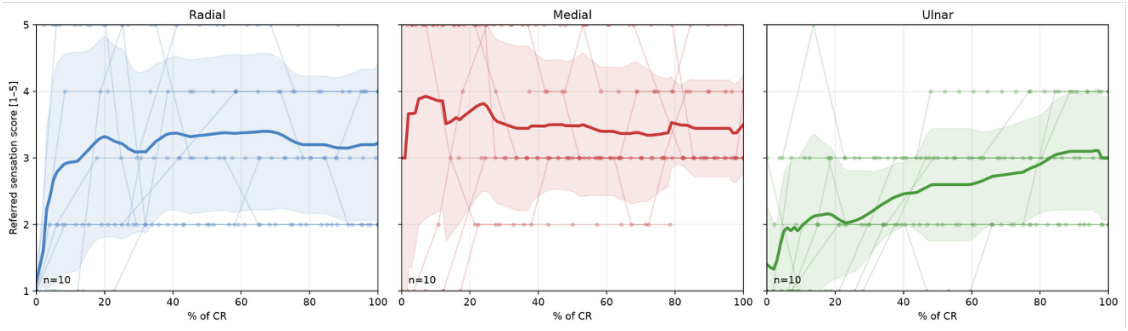


Figure 48: Group level referred sensation balance across the normalized comfort range.

An additional comparison tested whether an alternative electrode arrangement improved the medial percept. In the standard medial wrist configuration, both electrodes were placed on the volar wrist. In the medial opposite configuration, the stimulating electrode remained on the volar wrist while the return electrode was moved to the dorsal side, as illustrated in Figure 49.

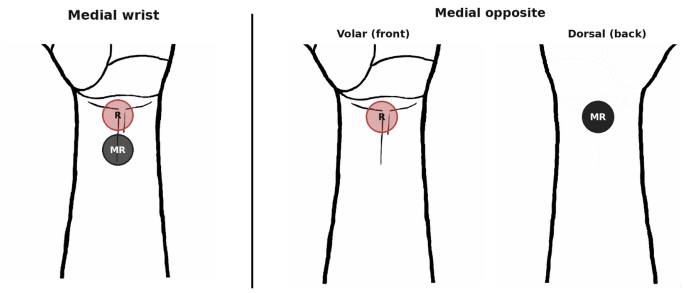


Figure 49: Standard medial wrist and medial opposite electrode configurations.

The opposite configuration did not produce a robust improvement in referred sensation balance. Although the group mean increased slightly, individual responses were highly variable. The projected fields nevertheless became more focused, particularly over the central fingers, suggesting a possible advantage when spatial selectivity is prioritized. This differs from the clearer reduction of local sensations reported by Pena et al. with a ventral-stimulation/dorsal-return arrangement [30]. The comparison is not direct, however: their channel-hopping system used independent current sources and precisely interleaved pulses, whereas the present experiment changed only the return-electrode position in a single-stimulator setup. Because the observed localization effect was not accompanied by a consistent improvement in the local-to-referred balance, the standard medial wrist configuration was retained for the subsequent experiments.

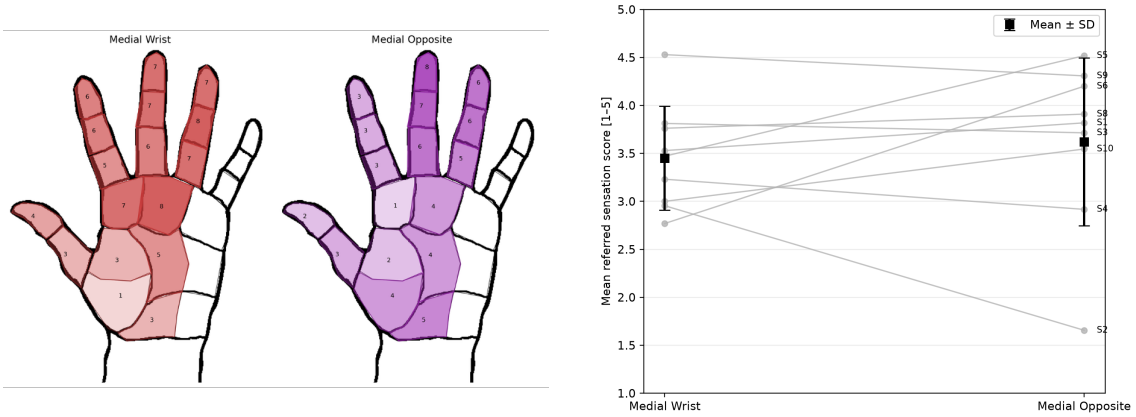


Figure 50: Projected fields and referred sensation balance for the two medial configurations.

5.2 Amplitude versus frequency discrimination on an ACR scale

5.2.1 Linear ACR analysis

The second experiment compared how finely subjects could discriminate changes produced through current amplitude or pulse frequency when both were expressed on the same linear ACR scale. The two modulation conditions used the same reference stimulus and the same relative ACR test levels, so differences in the fitted psychometric curves reflect how reliably each stimulation variable conveyed a change in perceived intensity.

Figure 51 compares the empirical 2AFC probabilities and logistic fits for one representative subject (A) and at group level (B). In both cases, the amplitude-modulation curve is steeper than the frequency-modulation curve, producing a sharper transition around the reference and indicating finer discrimination of amplitude-driven than frequency-driven ACR changes.

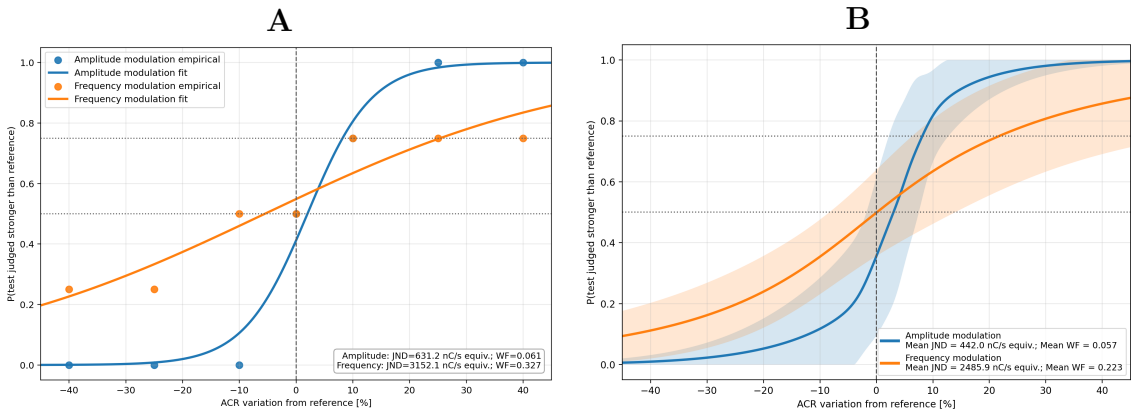


Figure 51: Linear ACR psychometric curves for a representative subject (A) and at group level (B).

The Weber fraction, defined in Section 4.3 as the JND normalized by the reference ACR, provides a relative measure of discrimination performance. As shown in Figure 52, the amplitude condition had a lower mean Weber fraction than the frequency condition. Subjects therefore required a smaller relative change when the ACR variation was produced through amplitude. This supports using amplitude as the primary control variable for perceived intensity or spatial recruitment, while frequency may be better suited to temporal quality and event dynamics. The result is consistent with implanted-nerve studies in which pulse amplitude or charge provided a direct and finely graded control of perceived intensity [28, 33].

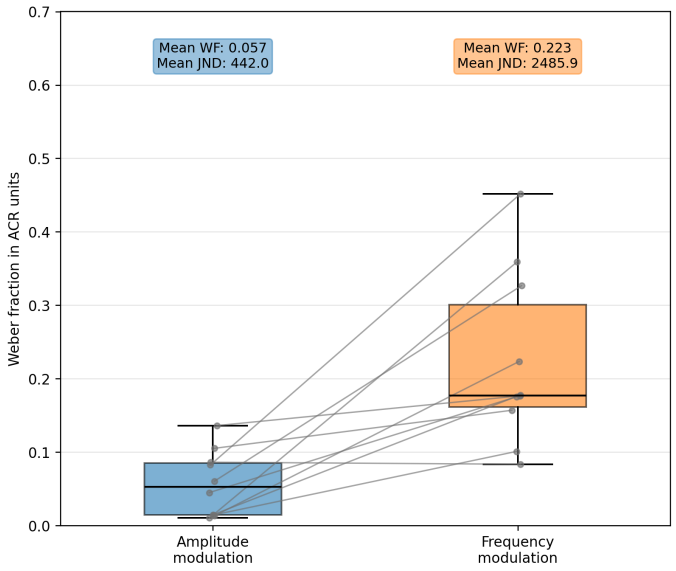


Figure 52: Weber fractions for amplitude and frequency modulation on the linear ACR scale.

5.2.2 Adapted logarithmic ACR scale

The linear ACR analysis indicated that frequency modulation produced a less steep psychometric curve and a higher Weber fraction than amplitude modulation. This suggests that, when frequency is treated as a linear multiplicative term in the ACR definition, frequency changes appear less finely discriminable than amplitude changes. A possible explanation is that the perceptual contribution of pulse frequency is not linear. In the work of Bucciarelli et al., this issue was addressed by introducing an adapted activation charge rate formulation in which the frequency contribution is compressed logarithmically, consistent with the general psychophysical observation that perceived intensity often scales nonlinearly with physical stimulus magnitude [29]. Following this idea, an adapted ACR value was computed as

$$\mathrm{ACR}_T = (I - I_{\min})\ \mathrm{PW}\ \log_2(f).$$

This transformation keeps the amplitude term linear but reduces the effective contribution of frequency increases. It was therefore used here as an exploratory analysis to test whether amplitude and frequency modulation become more comparable when frequency is represented on a logarithmic scale.

The transformation does not change the subject responses: the same 2AFC probabilities are used, but the x-axis coordinate associated with each frequency-modulation level is recomputed. For amplitude modulation, the frequency is fixed at f_{ref}, therefore, after normalization by the reference value, the logarithmic frequency term is constant and cancels out. As a consequence, the amplitude-modulation data are unchanged when represented as normalized variation from the reference.

For frequency modulation, current amplitude is fixed at I_{ref}, and the adapted normalized ratio becomes

$$\frac{\mathrm{ACR}_{T,\mathrm{test}}}{\mathrm{ACR}_{T,\mathrm{ref}}} = \frac{\log_2(f_{\mathrm{test}})}{\log_2(f_{\mathrm{ref}})}.$$

If the test frequency is generated from the nominal ACR factor k, so that $f_{\mathrm{test}} = k f_{\mathrm{ref}}$, then

$$\frac{\mathrm{ACR}_{T,\mathrm{test}}}{\mathrm{ACR}_{T,\mathrm{ref}}} = 1 + \frac{\log_2(k)}{\log_2(f_{\mathrm{ref}})}.$$

Thus, the frequency-modulation points are compressed around the reference on the adapted ACR axis. In practical terms, the probabilities remain those measured experimentally, but the horizontal spacing between frequency levels becomes smaller because equal linear changes in frequency do not correspond to equal changes on the logarithmic scale.

The same logistic fitting procedure described above was then repeated using the adapted ACR axis. For visual comparison with the original ACR plots, the fitted

curves were evaluated over the same full x-axis range. Therefore, especially for frequency modulation, the tails of the fitted ACRT curves outside the actually tested ACRT range should be interpreted as extrapolated portions of the logistic fit rather than as directly sampled experimental data.

Figure 53 shows the same representative subject (A) and the group result (B) after applying the adapted ACR transformation. The empirical response probabilities are unchanged, but the frequency-modulation points are shifted on the x-axis according to the logarithmic frequency scaling.

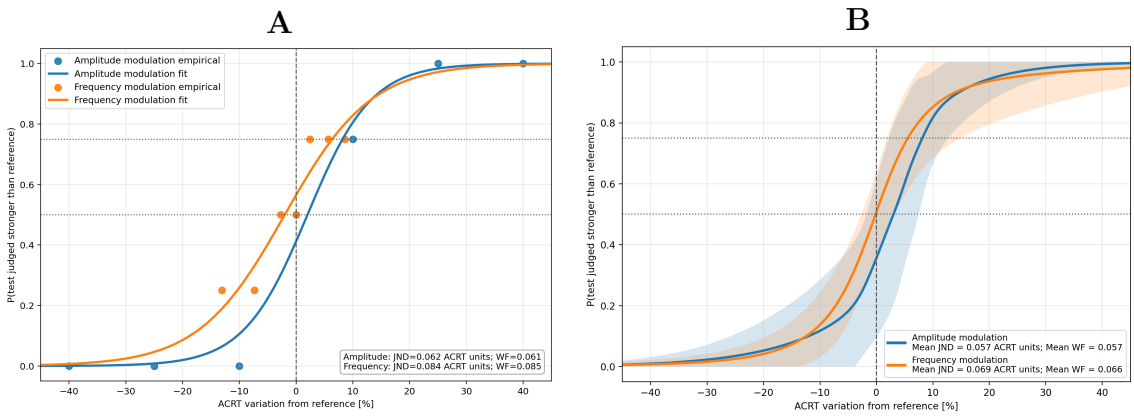


Figure 53: Adapted logarithmic ACR psychometric curves for a representative subject (A) and at group level (B).

Both representations show the same effect: the logarithmic transformation compresses the frequency-modulation points toward the reference, bringing the frequency curve closer to the amplitude curve. Their agreement supports a compressed, rather than purely linear, contribution of frequency to perceived intensity.

Figure 54 compares the Weber fractions after applying the adapted ACR transformation. In contrast to the linear ACR result, the mean Weber fractions for amplitude and frequency modulation are much closer, indicating that the logarithmic ACR representation reduces the apparent discrepancy between the two modulation variables.

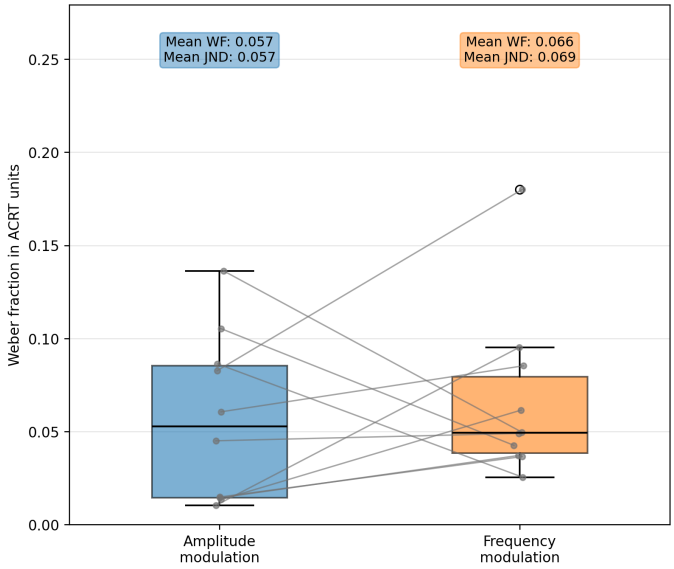


Figure 54: Weber fractions on the adapted logarithmic ACR scale.

Taken together, the linear ACR and adapted ACR analyses suggest that amplitude modulation provides a more direct and reliable way to control perceived intensity linearly, whereas frequency modulation has a non-linear perceptual contribution. At the same time, the adapted ACR analysis suggests that a logarithmic frequency scaling may be more appropriate when amplitude and frequency need to be compared on a shared perceptual activation scale.

5.3 Frequency characterization

The third experiment examined the perceptual role of frequency directly. Stimulation amplitude was fixed at 50% of the medial-channel comfort range, while pulse frequency varied from 10 to 300 Hz. Subjects rated both perceived intensity and temporal quality, allowing the effect of frequency on these two perceptual dimensions to be evaluated separately.

Perceived intensity increased steeply at low and intermediate frequencies and then tended to saturate at higher frequencies. As shown in Figure 55, the same group data become closer to a linear relationship when plotted against $\log_2(f)$. This independently confirms the result of the adapted ACR analysis: frequency contributes to perceived intensity, but its effect is compressed over the tested range rather than proportional to frequency itself.

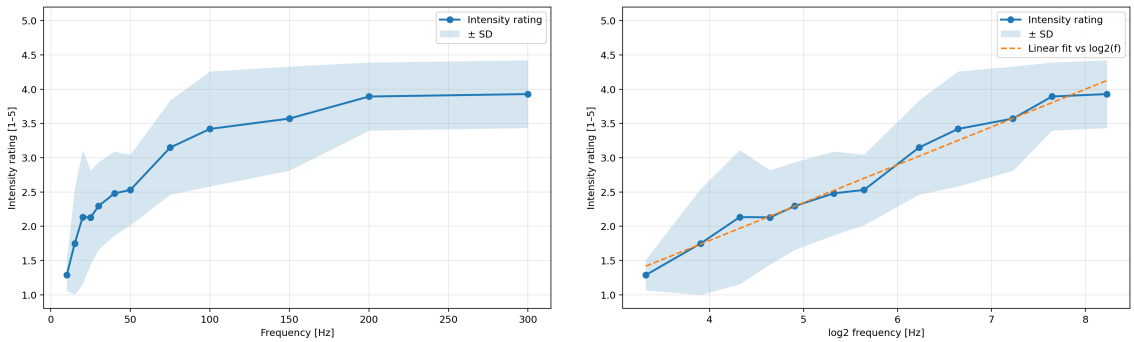


Figure 55: Group intensity ratings on linear-frequency and $\log_2$-frequency axes.

To quantify this difference, a linear-frequency model and a $\log_2$-frequency model were fitted separately for each subject. The representative example in Figure 56A shows that the logarithmic model follows the measured ratings more closely, with lower RMSE and higher R^2. At group level, the $\log_2$ model produced a higher R^2 in 8 of 10 subjects. The mean improvement was $\Delta R^2 = +0.126$ and was significant in the paired Wilcoxon signed-rank test ($p = 0.00684$; Figure 56B). The non-parametric paired test was used because the two model fits were evaluated in the same participants and the sample size was limited. These results do not establish an exact logarithmic perceptual law, but show that log-scaled frequency described this dataset better than linear frequency.

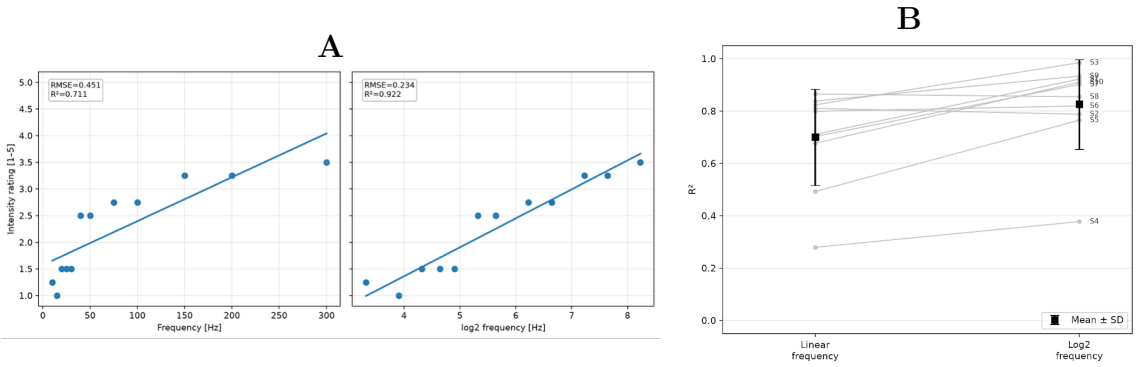


Figure 56: Linear and $\log_2$ frequency models for a representative subject (A) and paired R^2 comparison across subjects (B).

The category responses show that frequency also changed the temporal quality of the percept. Low frequencies were mainly described as tapping, with individual pulses perceived as intermittent events. Vibrating responses became more frequent at intermediate frequencies, while fused, continuous percepts were more likely at high frequencies (Figure 57). This frequency-dependent change in perceptual quality agrees with observations from peripheral nerve stimulation, where frequency shaped qualitative attributes in addition to intensity [31]. Across-subject variability was lower at the lower frequencies, where the transition from tapping

to vibrating was comparatively clear, and increased at higher frequencies, where the distinction between vibrating and fused percepts was less consistent.

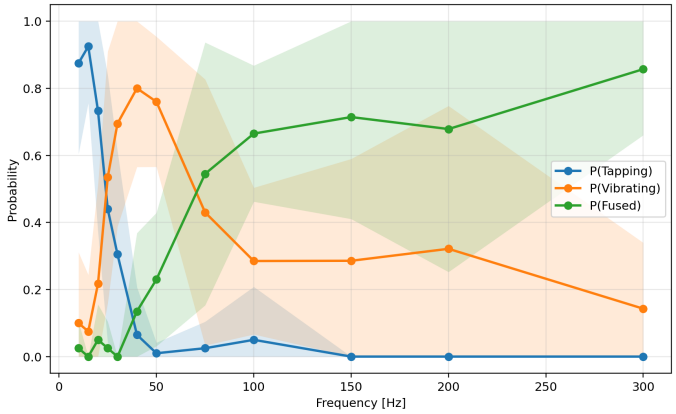


Figure 57: Group probabilities of tapping, vibrating, and fused percepts across frequency.

A representative subject is shown in Figure 58. Two transition frequencies were estimated: tapping to non-intermittent perception, where the combined vibrating and fused probability reached 50%, and transition to fused perception, where the fused probability reached 50%. Values between tested frequencies were obtained by linear interpolation.

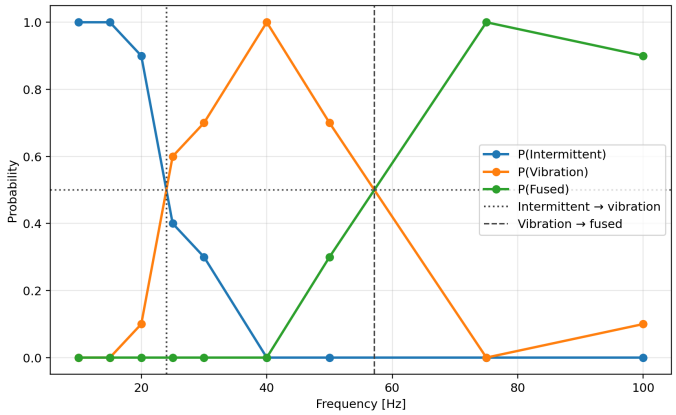


Figure 58: Representative subject category probability curves with estimated transition frequencies. The dotted horizontal line marks the 50% criterion.

Figure 59 shows that these transition frequencies varied substantially across subjects. As expected from the group probability curves, the tapping-to-non-intermittent transition was comparatively consistent, whereas the transition to fused perception occurred at higher frequencies and showed much greater variability. The aim was therefore not to identify universal frequency thresholds,

but to demonstrate a repeatable qualitative progression whose exact boundaries remain subject specific.

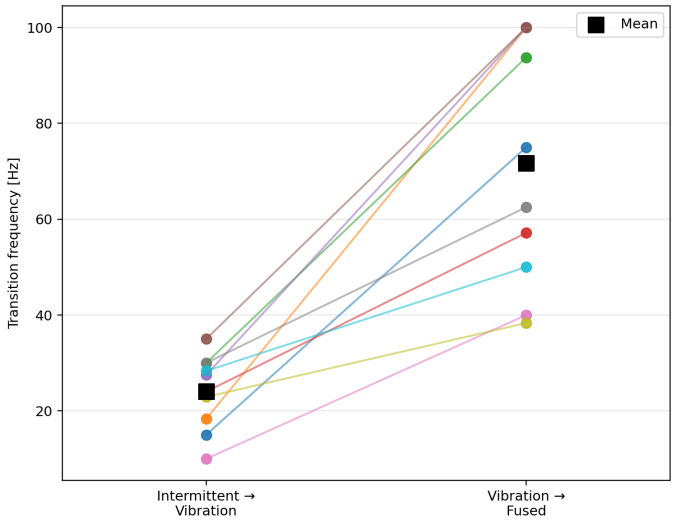


Figure 59: Subject level transition frequencies between perceptual categories. Each coloured line corresponds to one subject, and black squares indicate the group mean.

Taken together, the results distinguish the roles of the two real-time stimulation variables. Amplitude provided the finer and more direct control of intensity and spatial recruitment, whereas frequency influenced intensity non-linearly and more clearly controlled the temporal character of the percept. This supports using frequency to encode contact dynamics, transitions, and vibration-like qualities rather than treating it as a linear substitute for amplitude.

5.4 Overlapping train motion discrimination

The fourth experiment tested whether the onset delay between two overlapping stimulation trains could evoke a directional, motion-like percept across the medial and ulnar channels. The group mean clarity rating increased rapidly from 100 ms and reached its highest values at intermediate delays before decreasing at the longest delays (Figure 60). Very short delays therefore tended to produce simultaneous or unclear sensations, whereas very long delays were more likely to be perceived as two separate sequential events rather than as one continuous movement.

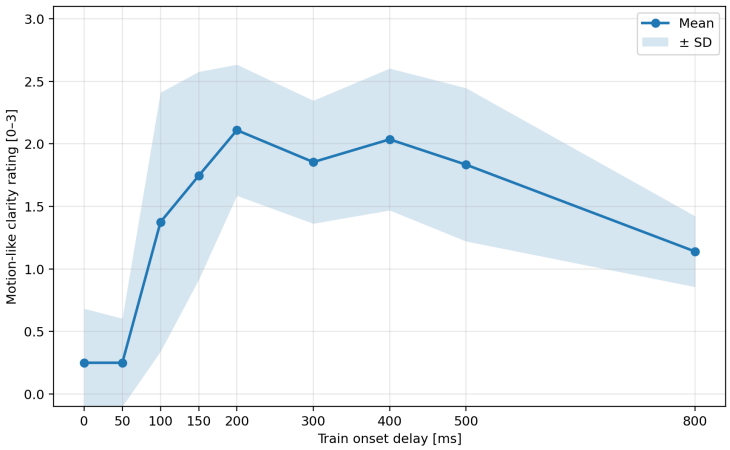


Figure 60: Mean across subjects of the motion like clarity rating as a function of train onset delay. Ratings range from 0 to 3, where 0 corresponds to no clear direction or simultaneous perception, and 3 corresponds to a clear motion like percept.

Figure 61 summarizes directional discrimination after collapsing the two directions. The probability of a correct response increased with delay, reaching high values from approximately 200–300 ms onward. The left panel reports the subject level d_{75}, defined in the Methods section as the onset delay at which the fitted probability of a correct directional response reaches 75%. Across subjects, the mean threshold was around 150 ms, although the individual points show visible inter-subject variability.

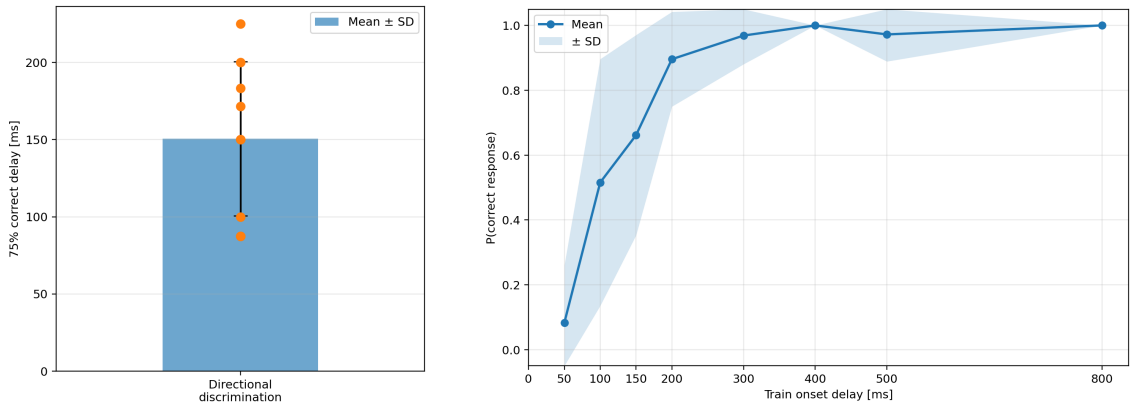


Figure 61: Directional discrimination summary. Left: subject level delays at which the 75% correct criterion was reached, with group mean and standard deviation. Right: group mean probability of correct directional response as a function of train onset delay, collapsed across A to B and B to A directions.

The two directions were then inspected separately to check for possible asymmetries between medial to ulnar and ulnar to medial stimulation. Figure 62 shows the direction specific probability curves. Both directions reached high

accuracy at intermediate and long delays. At shorter delays, variability was larger, indicating that some subjects could already identify the onset order while others still reported simultaneous or unclear direction.

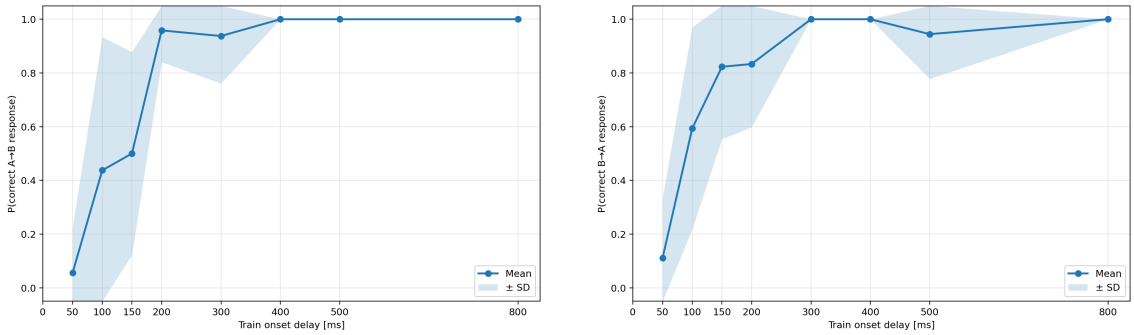


Figure 62: Direction specific probability of correct responses as a function of train onset delay. Left: A to B direction, corresponding to medial to ulnar activation. Right: B to A direction, corresponding to ulnar to medial activation.

To further test whether one direction required systematically longer delays than the other, the 75% discrimination delays obtained for A to B and B to A were compared within subjects using a paired Wilcoxon signed rank test. This non parametric paired test was chosen because the two thresholds were obtained from the same subjects and because the estimates were based on a limited number of discrete tested delays. No significant difference was observed between the two directions ($p = 0.094$), suggesting that directional discrimination was not strongly biased toward one stimulation order in this preliminary dataset.

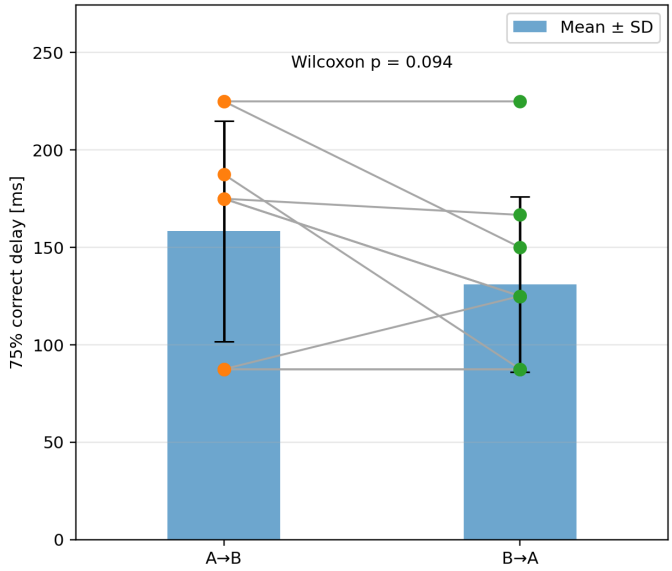


Figure 63: Paired comparison of the 75% directional discrimination delay for A to B and B to A stimulation. Grey lines connect the two values from the same subject. The p value was obtained using a paired Wilcoxon signed rank test.

Overall, this experiment shows that long delays can make the activation order easier to identify, but they may also reduce the continuity of the moving percept. Conversely, very short delays tend to produce fusion or simultaneity. The most useful range for future motion like encoding therefore appears to be an intermediate delay range, where directional discriminability and motion like clarity are both relatively high. This interpretation follows the classical apparent-haptic-movement dependence on stimulus duration and onset interval and is consistent with its later application to tTENS across projected hand locations [32, 16]. In the present work, only one burst duration was tested, so these results should be considered a preliminary characterization rather than a complete model of apparent movement.

5.5 Modulation strategy identification and naturalness

The final experiment evaluated how effectively the five encoding strategies conveyed the six simulated tactile scenarios. In no-visual trials, functional performance was measured from the scenario identified by the subject, while naturalness ratings were used to compare the perceptual quality of the feedback with and without visual context.

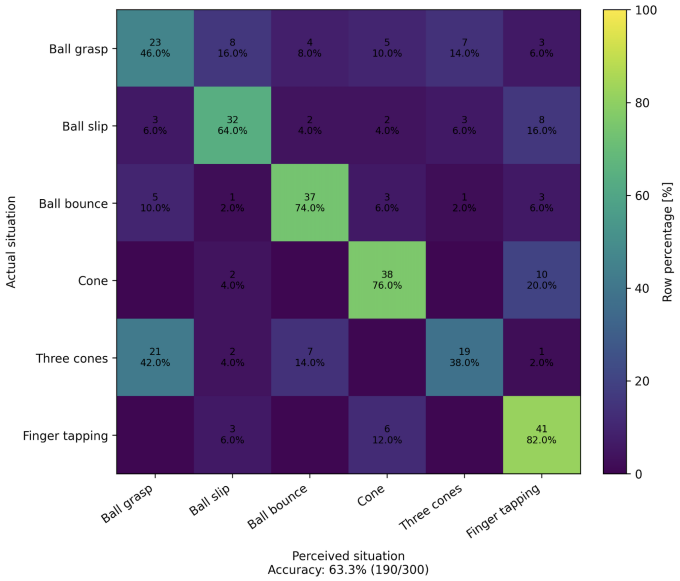


Figure 64: Overall no-visual identification confusion matrix across subjects and encoding strategies.

The overall identification accuracy was approximately 63%. This intermediate performance was appropriate for comparing the encoders: the task was neither at ceiling nor dominated by failures, leaving sufficient room for strategy-dependent differences to emerge. The overall confusion matrix in Figure 64 shows that ball grasp and three cones were frequently confused, plausibly because compression into three stimulation channels can make their profiles similar despite differences in the spatial extent of contact. Other recurrent confusions involved cone and finger tapping, which both contain localized intermittent contacts, and finger tapping and ball slip, which can both generate sequential channel activation.

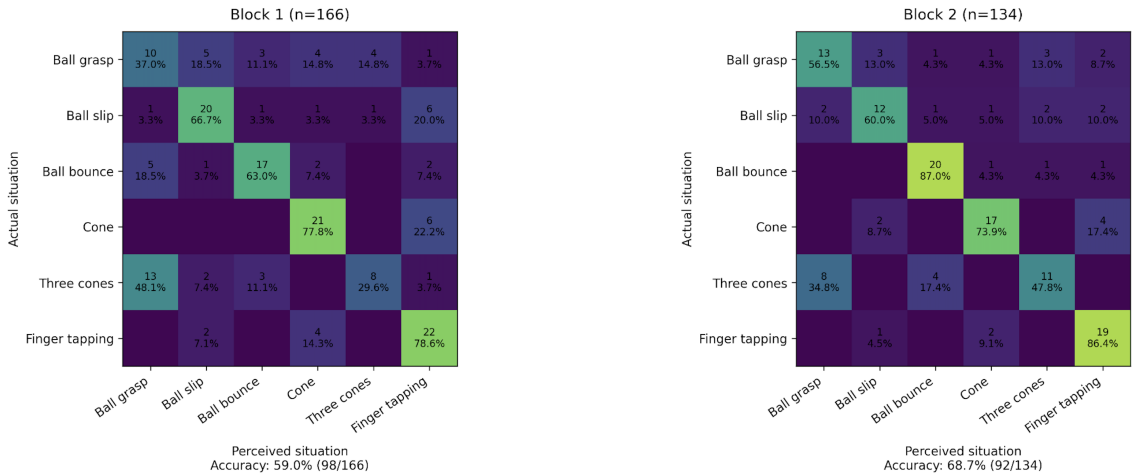


Figure 65: No-visual identification confusion matrices for the first and second experimental blocks.

The 60 trials for each subject were divided in two blocks of 30 by a break. Identification accuracy increased by approximately 10 percentage points from the first to the second block (Figure 65). This suggests that participants progressively learned the association between stimulation patterns and the available scenarios. Because visual and no-visual trials were randomly interleaved, the visual trials provided an explicit pairing between each simulated interaction and its stimulation pattern, thereby supporting learning that could subsequently improve no-visual identification. Performance therefore reflects both the information carried by an encoding strategy and the familiarity acquired during the session.

When the results are separated by encoding strategy (Figure 66), SA/RA LIF shows the clearest diagonal and reaches an identification accuracy of approximately 83%, substantially above the overall accuracy of 63%. More generally, strategies in which pressure controlled stimulation amplitude performed better than those relying primarily on frequency modulation. This is consistent with the second experiment, where amplitude provided a more effective representation of perceived intensity, while frequency was better suited to conveying temporal information. The SA/RA LIF strategy combined these two roles: amplitude followed pressure directly, while frequency represented the sustained and transient components extracted by the SA and RA pathways.

The more detailed and computationally demanding TouchMime strategies did not provide a corresponding identification advantage. TouchMime predicts structured activity across afferent populations, whereas tTENS recruits mixed populations of fibres with limited selectivity because current spreads through the electrode–skin interface and surrounding tissues. Part of the information generated by the model may therefore be lost when converted into transcutaneous stimulation. This contrasts with invasive intraneural interfaces, where greater fibre selectivity allows biomimetic temporal patterns to improve naturalness and functional performance [34]. Encoding strategies should therefore be evaluated across the complete sensor–interface–user chain rather than from the biological detail of the model alone.

Naturalness was analysed only for correctly identified no-visual trials. If a participant identified a different scenario, the rating referred to the perceived scenario rather than to the one actually presented and could not be compared meaningfully with the corresponding visual trial. Restricting the analysis to correct responses therefore ensures that visual and no-visual ratings refer to the same tactile situation. In all naturalness plots, bars represent the mean across subjects and error bars the inter-subject standard deviation. Statistical significance is indicated by one asterisk for $p < 0.05$, two asterisks for $p < 0.01$, and “ns” for a non-significant difference.

When all scenarios and strategies were pooled, visual feedback produced significantly higher naturalness ratings than the no-visual condition (Figure 67; two-sided Wilcoxon signed-rank test, $p < 0.01$). The visual condition also showed

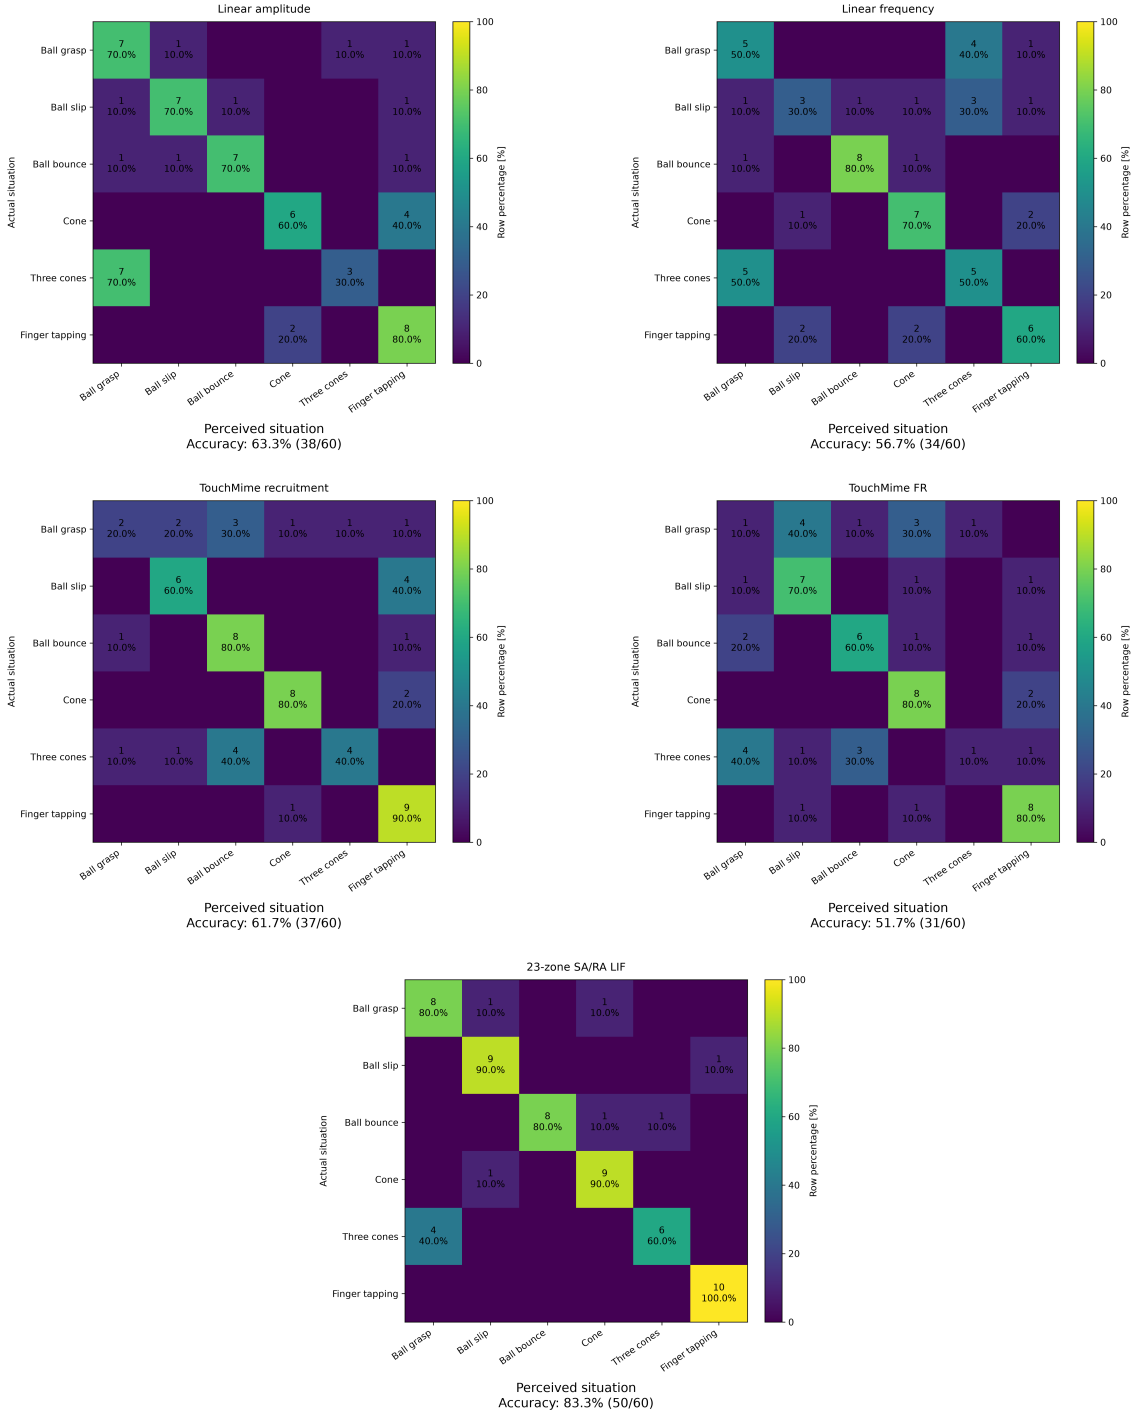


Figure 66: No-visual identification confusion matrices separated by encoding strategy.

lower variability, suggesting that visual context made the stimulation patterns more consistent to interpret across participants. The following plots report the same comparison separately by scenario (Figure 68) and by encoding strategy (Figure 69). Visual feedback produced a higher mean in every comparison, although not every difference was statistically significant. Across strategies, linear amplitude and SA/RA LIF remained among the most natural encodings when the scenario was correctly identified.

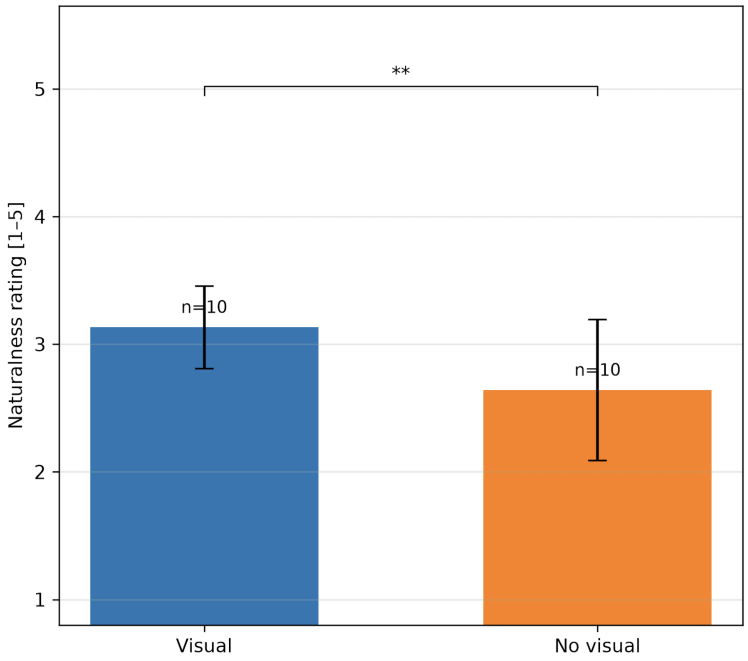


Figure 67: Overall correct-trial naturalness ratings with all scenarios and encoding strategies pooled. Bars show the mean across subjects and error bars show the inter-subject standard deviation.

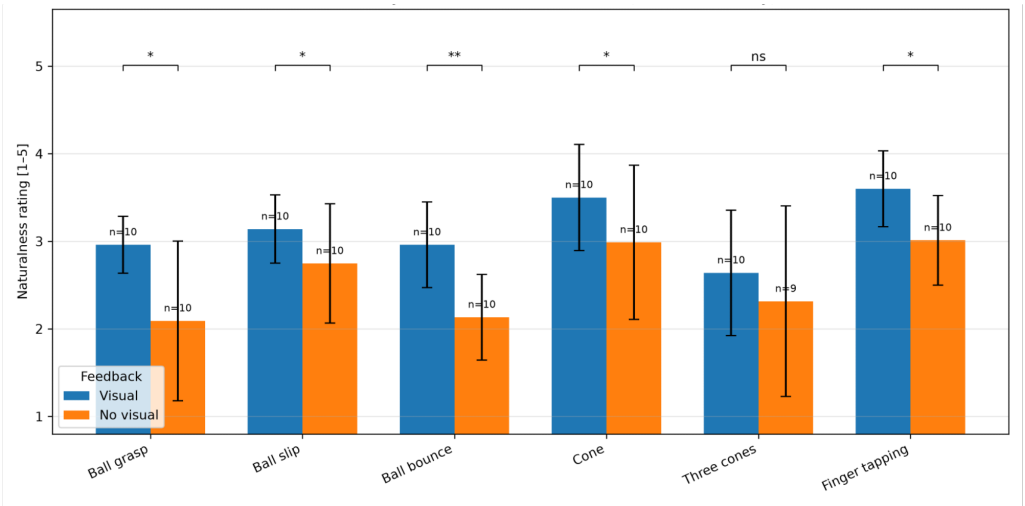


Figure 68: Naturalness by scenario after excluding incorrectly identified no-visual trials.

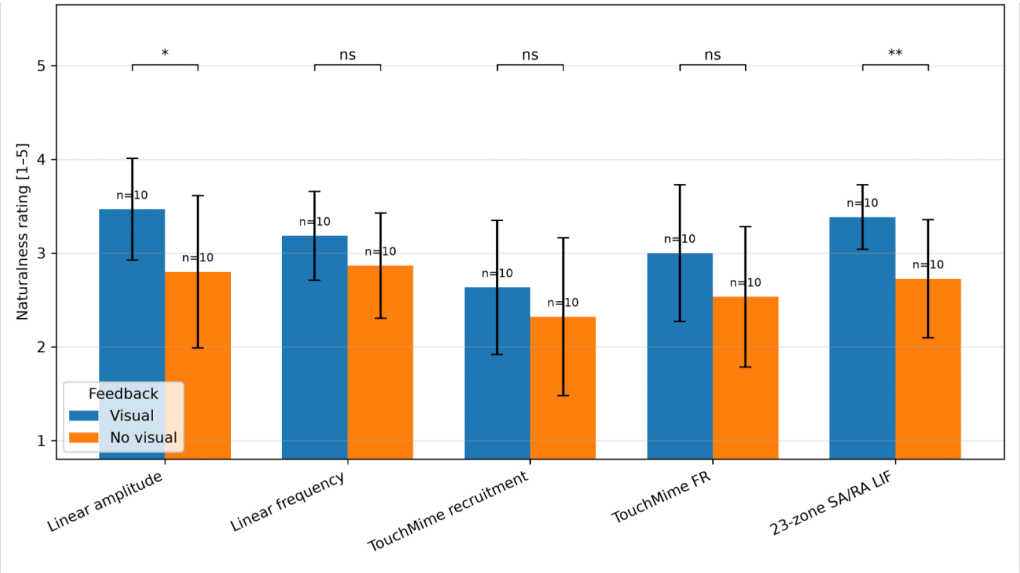


Figure 69: Naturalness by encoding strategy after excluding incorrectly identified no-visual trials.

6 Conclusion

6.1 Main contributions

This thesis developed and evaluated a complete pipeline for delivering multichannel tactile feedback through targeted transcutaneous electrical nerve stimulation. The work combined hardware and software integration, subject-specific calibration, characterization of the principal stimulation variables, evaluation of inter-channel temporal patterns, and comparison of five tactile encoding strategies in the same cohort of ten healthy participants. Its main high-level contribution is therefore not a single stimulation mapping, but an experimentally grounded account of which modulation and encoding principles remain effective after rich tactile information has been compressed into the limited control space of a non-invasive interface.

The calibration experiments showed that comfort ranges and projected fields were strongly dependent on both participant and stimulation channel. A nominally identical current could produce different perceived intensities, local sensations, and referred hand regions across subjects. Subject-specific mapping and threshold estimation must therefore be treated as integral parts of the feedback system rather than as preliminary technical adjustments. Comparisons between encoding strategies would otherwise be confounded by differences in accessibility and comfort that are unrelated to the encoding itself.

Amplitude and frequency were not perceptually interchangeable variables. Amplitude provided a comparatively direct and finely graded means of controlling perceived intensity and recruitment within the individual comfort range. Frequency also influenced intensity, but its perceptual effect was compressed and closer to a logarithmic than a linear progression. More importantly, frequency changed the temporal quality of the sensation, with the dominant percept progressing from tapping to vibration and then toward a fused sensation. These results support using amplitude as the primary continuous intensity variable, while frequency can provide an additional dimension for representing temporal quality or events when its nonlinear perceptual scale is explicitly considered.

The motion experiments further showed that the relative timing between channels carries usable information. For 500 ms stimulation trains, delays of approximately 150–300 ms produced the clearest apparent movement. Shorter delays tended to be integrated as simultaneous stimulation, whereas longer delays increasingly produced two separate events. This identifies an operating interval for encoding directional or slip-like changes through sequential stimulation, while also demonstrating that multichannel feedback cannot be designed solely from the independent properties of each channel.

The comparison of encoding strategies provided the clearest evidence for the overall design principle emerging from this work. In the no-visual scenario-identification

task, the SA/RA LIF strategy reached an accuracy of approximately 83%, compared with an overall accuracy of approximately 63%, and it remained among the most natural strategies together with linear amplitude modulation. Its advantage was obtained with a relatively simple biological abstraction: sustained pressure and transient changes were separated before being converted into stimulation commands. The more detailed TouchMime strategies did not yield a corresponding perceptual benefit.

This result should not be interpreted as evidence that biological realism is generally ineffective. Rather, it indicates that the appropriate level of realism depends on the stimulation interface. Fibre-level models predict structured activity across distinct afferent populations, but tTENS does not reproduce that selectivity. Current must cross the electrode–skin interface and surrounding tissues, where skin impedance, electrode geometry, nerve depth, and current spread transform the applied waveform and recruit mixed populations of superficial and deeper fibres. The resulting neural activation can therefore differ substantially from the simulated afferent pattern. Detailed fibre-level encoding may be better exploited by selective implanted interfaces, whereas non-invasive stimulation may benefit more from simple biomimetic principles adapted to its limited number of channels, participant-specific comfort ranges, spatial uncertainty, and hardware bandwidth. Within these constraints, the present results support a hybrid design philosophy: preserve the biologically important distinction between sustained and transient information, but optimize its final stimulation representation for what the user can reliably perceive.

6.2 Limitations

The conclusions are limited to the tested healthy cohort and experimental setup. The sample comprised ten participants, and the final strategy comparison was exploratory, with a limited number of repetitions for each scenario and strategy. The reported identification accuracies therefore describe the present cohort and protocol rather than establishing population-level superiority. In addition, naturalness and scenario identification were subjective psychophysical outcomes, and the study did not assess learning over extended use or stability across multiple sessions. The hardware also constrained the conclusions that can be drawn about multichannel encoding. Three projected territories were stimulated through a single current source and a multiplexer, so the channels were addressed sequentially rather than by independent simultaneous stimulators. The resulting total-frequency limit and shared timing architecture may have affected strategies that generated dense or rapidly changing commands. Moreover, the tactile inputs were simulated pressure profiles rather than signals acquired online from sensors integrated into a functional prosthesis. The complete closed loop therefore remains to be validated during active object manipulation.

Finally, results obtained in healthy participants do not replace validation in people with upper-limb amputation. Residual nerve anatomy, nerve depth, tissue changes, reinnervation, and phantom-hand representation may alter both accessibility and discriminability of the evoked sensations. These factors may either reduce performance or create more distinct stimulation sites than those observed in intact limbs, and this cannot be inferred from the present data alone.

6.3 Future possible directions

The most direct scientific extension is to develop the SA/RA LIF strategy further and determine whether its performance can be retained while progressively increasing its biological plausibility. The present model intentionally represents only the broad distinction between slowly and rapidly adapting responses. Sustained and transient components are extracted using low-pass and high-pass filters with the same empirically selected cutoff frequency, without distinguishing type-I from type-II afferents or deriving the temporal parameters from physiological measurements. This made the encoder an intermediate solution between the linear strategies and the more detailed TouchMime models, but it also leaves several mechanistic hypotheses untested.

A future receptor-specific implementation could provide separate SA-I, SA-II, RA-I, and RA-II pathways. The slowly adapting branches could use different low-frequency dynamics and adaptive spiking neurons so that their sustained firing and spike-frequency adaptation depend on receptor class. Rapidly adapting branches could use high-pass or band-pass stages and change-sensitive event encoders, including asynchronous delta modulation, to represent contact onset, release, slip, and vibration. The parameters should first be fitted to physiological afferent responses and then retuned, if necessary, for the bandwidth and selectivity of the stimulation interface. Existing neuromorphic tactile circuits demonstrate several feasible implementations: capacitive sensors with circuits that retain activity during sustained pressure can reproduce SA-like behaviour; event-driven POSFET taxels, which combine a pressure-sensitive transistor with local encoding electronics, can generate separate sustained and transient outputs; and CMOS circuits can reproduce simplified SA-I and RA-I afferent responses directly in low-power hardware [20, 21, 22].

This development also creates a route toward a low-power hardware implementation. Receptor-specific filters, adaptation mechanisms, and spike generators could first be reproduced and validated at circuit level using the tactile profiles developed in this thesis. Analog or mixed-signal implementations could subsequently be interfaced directly with tactile sensors, reducing the need to sample and process every pressure value in software. Such event-driven architectures are particularly relevant for wearable or prosthetic systems because computation and communication occur primarily when tactile information changes [23].

A more radical architecture would preserve the spike representation through the stimulation stage rather than averaging it and converting it back into a continuous frequency command. Each selected encoder event could trigger a charge-controlled stimulation pulse or a short, safety-constrained pulse packet through a dedicated output stage. This concept is technically plausible: neuromorphic tactile sensor spikes have previously been used as triggers for elementary electrical stimuli delivered through a multichannel skin interface in an animal preparation [24]. Biomimetic temporal stimulation patterns have also produced more natural neural dynamics and functional benefits with implanted peripheral interfaces [25]. However, a literal one-spike–one-pulse mapping should not be assumed to transfer directly to human tTENS. It would require charge limiting, refractory and maximum-rate constraints, channel arbitration, isolation, and psychophysical validation because transcutaneous current spread recruits populations of fibres rather than reproducing an individual biological spike. A programmable isolated current driver would therefore be more appropriate than an unspecified power amplifier.

Future hardware should also investigate alternative electrodes and return-electrode arrangements. The medial opposite configuration produced more spatially focused projected fields in some participants, particularly over the central fingers, although it did not robustly improve the balance between local and referred sensations. This observation motivates a systematic comparison of electrode size, shape, spacing, material, and return location. If a more selective opposite configuration can be made stable across subjects, it may support more than the three radial, medial, and ulnar channels used here. Increasing the number of useful channels would, however, require hardware with genuinely independent stimulation outputs. A multichannel stimulator capable of simultaneous or precisely interleaved pulses would remove the total-frequency bottleneck imposed by routing a single DS8R through the D188 multiplexer and would allow channel-specific amplitudes and timings to be controlled without reducing the available per-channel rate.

Another useful analysis is to quantify how much scenario information is lost at each stage of the encoding pipeline. A classifier can first test the separability of the original 23-zone pressure profiles and then repeat the analysis after spatial pooling and after conversion into the amplitude and frequency commands of each strategy. Confusion matrices and cross-validated classification accuracy could indicate whether errors such as confusion between ball grasp and three cones originate from similar encoded signals or from limitations in human perception and learning.

Finally, the next experimental study should reproduce the full calibration and strategy-comparison pipeline in participants with upper-limb amputation. In addition to identification accuracy and naturalness, it should measure stability across sessions, phantom-hand localization, functional prosthesis control, cognitive load, and robustness to changes in posture and electrode placement. The central

question is not whether a more detailed model is biologically convincing in isolation, but whether the complete sensor–encoder–stimulator–user loop communicates more useful information under realistic operating constraints.

References

[1] R. S. Johansson and A. B. Vallbo, "Tactile sensibility in the human hand: relative and absolute densities of four types of mechanoreceptive units in glabrous skin," *The Journal of Physiology*, vol. 286, pp. 283–300, 1979, doi: 10.1113/jphysiol.1979.sp012619.

[2] D. Deflorio, M. Di Luca, and A. M. Wing, "Skin and mechanoreceptor contribution to tactile input for perception: a review of simulation models," *Frontiers in Human Neuroscience*, vol. 16, p. 862344, 2022, doi: 10.3389/fnhum.2022.862344.

[3] B. P. Delhaye, K. H. Long, and S. J. Bensmaia, "Neural basis of touch and proprioception in primate cortex," *Comprehensive Physiology*, vol. 8, no. 4, pp. 1575–1602, 2018, doi: 10.1002/cphy.c170033.

[4] S. Standring, Ed., *Gray's Anatomy: The Anatomical Basis of Clinical Practice*, 42nd ed. Elsevier, 2021.

[5] 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. 738–750, 2022, doi: 10.1109/TNSRE.2022.3159189.

[6] P. Kourtesis, F. Argelaguet, S. Vizcay, M. Marchal, and C. Pacchierotti, "Electrotactile feedback applications for hand and arm interactions: a systematic review, meta-analysis, and future directions," *IEEE Transactions on Haptics*, vol. 15, no. 3, pp. 479–496, 2022, doi: 10.1109/TOH.2022.3189866.

[7] S. Raspopovic et al., "Restoring natural sensory feedback in real-time bidirectional hand prostheses," *Science Translational Medicine*, vol. 6, no. 222, p. 222ra19, 2014, doi: 10.1126/scitranslmed.3006820.

[8] J. A. George et al., "Biomimetic sensory feedback through peripheral nerve stimulation improves dexterous use of a bionic hand," *Science Robotics*, vol. 4, no. 32, p. eaax2352, 2019, doi: 10.1126/scirobotics.aax2352.

[9] 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, p. 35, 2015, doi: 10.1186/s12984-015-0022-8.

[10] M. Franceschi, L. Seminara, L. Pinna, S. Dosen, D. Farina, and M. Valle, "A system for electrotactile feedback using electronic skin and flexible matrix electrodes: experimental evaluation," *IEEE Transactions on Haptics*, vol. 10, no. 2, pp. 162–172, 2017, doi: 10.1109/TOH.2016.2618377.

[11] G. Chai, X. Sui, S. Li, L. He, and N. Lan, "Characterization of evoked tactile sensation in forearm amputees with transcutaneous electrical nerve stimulation," *Journal of Neural Engineering*, vol. 12, no. 6, p. 066002, 2015, doi: 10.1088/1741-2560/12/6/066002.

[12] E. D'Anna et al., "A somatotopic bidirectional hand prosthesis with transcutaneous electrical nerve stimulation based sensory feedback," *Scientific Reports*, vol. 7, p. 10930, 2017, doi: 10.1038/s41598-017-11306-w.

[13] L. E. Osborn et al., "Targeted transcutaneous electrical nerve stimulation for phantom limb sensory feedback," in *2017 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pp. 1–4, 2017, doi: 10.1109/BIOCAS.2017.8325200.

[14] H. Shin, Z. Watkins, H. H. Huang, Y. Zhu, and X. Hu, "Evoked haptic sensations in the hand via non-invasive proximal nerve stimulation," *Journal of Neural Engineering*, vol. 15, no. 4, p. 046005, 2018, doi: 10.1088/1741-2552/aabd5d.

[15] L. Vargas, G. Whitehouse, H. H. Huang, Y. Zhu, and X. Hu, "Evoked haptic sensation in the hand with concurrent non-invasive nerve stimulation," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 10, pp. 2761–2767, 2019, doi: 10.1109/TBME.2019.2895575.

[16] A. Scarpelli et al., "Evoking apparent moving sensation in the hand via transcutaneous electrical nerve stimulation," *Frontiers in Neuroscience*, vol. 14, p. 534, 2020, doi: 10.3389/fnins.2020.00534.

[17] M. Hao et al., "Restoring finger-specific sensory feedback for transradial amputees via non-invasive evoked tactile sensation," *IEEE Open Journal of Engineering in Medicine and Biology*, vol. 1, pp. 98–107, 2020, doi: 10.1109/OJEMB.2020.2981566.

[18] L. Vargas, H. H. Huang, Y. Zhu, and X. Hu, "Object recognition via evoked sensory feedback during control of a prosthetic hand," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 207–214, 2022, doi: 10.1109/LRA.2021.3122897.

[19] A. Scarpelli et al., "Eliciting force and slippage in upper limb amputees through transcutaneous electrical nerve stimulation (TENS)," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024, doi: 10.1109/TNSRE.2024.3443398.

[20] S. Caviglia, L. Pinna, M. Valle, and C. Bartolozzi, "An event-driven POSFET taxel for sustained and transient sensing," in *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 349–352, 2016, doi: 10.1109/ISCAS.2016.7527242.

[21] E. Rahiminejad, A. Parvizi-Fard, M. M. Iskarous, N. V. Thakor, and M. Amiri, "A biomimetic circuit for electronic skin with application in hand prosthesis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 2333–2344, 2021, doi: 10.1109/TNSRE.2021.3120446.

[22] E. Janotte, S. Bamford, O. Richter, M. Valle, and C. Bartolozzi, "Neuromorphic capacitive tactile sensors inspired by slowly adaptive mechanoreceptors," in *2022 20th IEEE Interregional NEWCAS Conference (NEWCAS)*, pp. 119–123, 2022, doi: 10.1109/NEWCAS52662.2022.9841988.

[23] E. Donati and G. Valle, "Neuromorphic hardware for somatosensory neuroprostheses," *Nature Communications*, vol. 15, p. 556, 2024, doi: 10.1038/s41467-024-44723-3.

[24] C. M. Oddo et al., "Artificial spatiotemporal touch inputs reveal complementary decoding in neocortical neurons," *Scientific Reports*, vol. 7, p. 45898, 2017, doi: 10.1038/srep45898.

[25] G. Valle et al., "Biomimetic computer-to-brain communication enhancing naturalistic touch sensations via peripheral nerve stimulation," *Nature Communications*, vol. 15, p. 1151, 2024, doi: 10.1038/s41467-024-45190-6.

[26] Digitimer Ltd., *DS8R Biphasic Constant Current Stimulator: Operator's Manual*, version 2.1. Available: `https://www.digitimer.com/product/life-science-research/stimulators/ds8r-biphasic-constant-current-stimulator/` (accessed Aug. 23, 2026).

[27] Digitimer Ltd., *D188 Remote Electrode Selector: Operator's Manual*. Available: `https://www.digitimer.com/wp-content/uploads/2019/09/D188-Remote-Electrode-Selector.pdf` (accessed Aug. 23, 2026).

[28] E. L. Graczyk, M. A. Schiefer, H. P. Saal, B. P. Delhaye, S. J. Bensmaia, and D. J. Tyler, "The neural basis of perceived intensity in natural and artificial touch," *Science Translational Medicine*, vol. 8, no. 362, p. 362ra142, 2016, doi: 10.1126/scitranslmed.aaf5187.

[29] V. Bucciarelli, N. Gozzi, N. Katic, G. Aiello, M. Razzoli, G. Valle, and S. Raspopovic, "Multiparametric non-linear TENS modulation to integrate intuitive sensory feedback," *Journal of Neural Engineering*, vol. 20, no. 3, p. 036026, 2023, doi: 10.1088/1741-2552/acd4e8.

[30] A. E. Pena, J. J. Abbas, and R. Jung, "Channel-hopping during surface electrical neurostimulation elicits selective, comfortable, distally referred sensations," *Journal of Neural Engineering*, vol. 18, p. 055004, 2021, doi: 10.1088/1741-2552/abf28c.

[31] E. L. Graczyk, B. P. Christie, Q. He, D. J. Tyler, and S. J. Bensmaia, "Frequency shapes the quality of tactile percepts evoked through electrical stimulation of the nerves," *The Journal of Neuroscience*, vol. 42, no. 10, pp. 2052–2064, 2022, doi: 10.1523/JNEUROSCI.1494-21.2021.

[32] C. E. Sherrick and R. Rogers, "Apparent haptic movement," *Perception & Psychophysics*, vol. 1, pp. 175–180, 1966, doi: 10.3758/BF03215780.

[33] G. Valle et al., "Comparison of linear frequency and amplitude modulation for intraneural sensory feedback in bidirectional hand prostheses," *Scientific Reports*, vol. 8, p. 16666, 2018, doi: 10.1038/s41598-018-34910-w.

[34] G. Valle et al., "Biomimetic intraneural sensory feedback enhances sensation naturalness, tactile sensitivity, and manual dexterity in a bidirectional prosthesis," *Neuron*, vol. 100, pp. 37–45.e7, 2018, doi: 10.1016/j.neuron.2018.08.033.