Ear EEG: Sensors and Systems for User-generic Neural Hearables

Ryan Kaveh

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2023-56
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2023/EECS-2023-56.html

May 1, 2023

Acknowledgement

This work is the culmination of many collaborations and technical discussions. Firstly, I would like to thank my dissertation committee, Professors Rikky Muller, Ana Claudia Arias, and Robert Knight. Professor Muller has been an invaluable mentor who has given me the freedom to think outside the box and develop quirky solutions to crazy problems. She has come to my defense countless times and I can always trust that she will be in my corner to back me up. In addition, Professors Arias and Knight have both provided amazing technical, material, and emotional support. All three of these mentors have enhanced my graduate career at UC Berkeley.

The portal won't let me add the rest of my acknowledgments so please read my thesis to see the names of the many collaborators.

Ear EEG: Sensors and Systems for User-generic Neural Hearables

by

Ryan C Kaveh

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Electrical Engineering and Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Rikky Muller, Chair
Professor Ana C. Arias
Professor Robert T. Knight

Summer 2022

The dissertation of Ryan C Kaveh, titled Ear EEG: Sensors and Systems for User-generic Neural Hearables, is approved:

Chair _____ Date _____

_____ Date _____

_____ Date _____

University of California, Berkeley

Ear EEG: Sensors and Systems for User-generic Neural Hearables

Abstract

Ear EEG: Sensors and Systems for User-generic Neural Hearables

by

Ryan C Kaveh

Doctor of Philosophy in Electrical Engineering and Computer Science

University of California, Berkeley

Professor Rikky Muller, Chair

In recent decades, wearable health monitors have grown from crude heart rate sensors to all-in-one devices that can track steps, motion, location, arrhythmia, electrocardiogram, blood oximetry, and calorie usage. More recent work has focused on developing wearables that provide non-invasive neural recording (electroencephalography - EEG) for focus, stress, and drowsiness monitoring. As wrist-worn wearables run out of features to add, head/ear worn EEG wearables offer new ways to provide users with helpful and actionable wellness information. The problem blocking widespread neural wearables is that existing devices are bulky, uncomfortable, require single-use wet electrodes, and often require training in everyday scenarios.

This thesis details an end-to-end design process for low-profile, dry-electrode neural recording hearables that can record EEG inside the ear. Starting from the EEG signal basics, this work will walk through the modeling, design, and verification of the three parts of an Ear EEG system: the electrodes, the neural recording system, and the downstream processing and classification software. Particular focus is placed on maximizing user comfort and the ease of manufacture without hurting system performance.

All of these topics will be based on constructing a practical in-ear EEG device based on multiple dry electrodes, a user-generic design, and a lightweight wireless interface for streaming data and device programming. Different earpiece manufacturing processes will be showcased for prototyping (using 3D printing and electroless plating) and production at scale (using vacuum forming and spray coating). The performance of this system will be evaluated with human subject trials that recorded spontaneous and evoked physiological signals, eye-blinks, alpha rhythm, auditory steady-state response (ASSR), and drowsiness detection.

To my family

My parents, Hedy and Sherry Kaveh, and my siblings, Anthony and Ariana Kaveh, have been my greatest single source of encouragement, motivation, humor, and happiness. I love them so much it genuinely hurts sometimes. I've known them for 27 years at this point - well everyone, but Ariana. I've only known her for 25 years, I guess. Regardless, these last 25 years have been the best years of my life and the common thread is my family.

Contents

List of Figures

List of Tables

Acknowledgments

This work is the culmination of many collaborations and technical discussions. Firstly, I would like to thank my dissertation committee, Professors Rikky Muller, Ana Claudia Arias, and Robert Knight. Professor Muller has been an invaluable mentor who has given me the freedom to think outside the box and develop quirky solutions to crazy problems. She has come to my defense countless times and I can always trust that she will be in my corner to back me up. In addition, Professors Arias and Knight have both provided amazing technical, material, and emotional support. All three of these mentors have enhanced my graduate career at UC Berkeley.

My labmates and collaborators have all helped push me and support me when everything happened to break at once. Justin Doong, we have worked together from the start and without you, none of this would have been possible. In fact, all of the work presented in this thesis is the result of collaborations so I would like to also thank Carolyn Schwendeman, Karthik Gopalan, Natalie Tetreault, Julian Maravilla, Aviral Pandey, Sina Faraji Alamouti, Andy Zhou, Cem Yalcin, Meraj Ghanbari, Fred Burghardt, Professor Michel Maharbiz, and Professor Miki Lustig.

I would also like to emphasize that this Ph.D. was not a solo endeavor. Clearly, I've had technical/physical help along the way from others inside the program. However, my friends outside the program have helped buoy me when it felt like I was drowning. Whether it be through walking Karma with Sneha, running, house-mating, climbing, gaming, wololo-ing, hiking, surfing, cycling, lifting, porch wine-ing, sitting, bae-ing, talking, kinda backpacking, dedar-ing, vegetating, laying, hugging, smiling, joking, working, or driving - my friends have been a truly constant source of motivation, peace, and happiness. I've looked up to all of my friends and I'm very grateful that they've been with me for so long.

Lastly, I would like to acknowledge my mentors and role models from outside my research: Dr. Anthony Kaveh (hah), Shirley Salanio, Professor Luke Theogarajan, Dr. Aaron Bluestone, Dr. Ian Morgan, Andy Sigears, and Albert Caruana. One way or another - for better or for worse - I have tried to emulate traits from all of these amazing people. Through high school and college, these mentors have inspired me and helped shape my worldview. I owe so much to all of them and without them, I would not have had such a curious spirit and peculiar work ethic.

Chapter 1

Introduction, EEG Signals, and Recording Systems

1.1 Introduction

Wearable devices (wearables) have revolutionized at-home wellness monitoring. Devices such as cell phones, Fitbit, the Apple Watch, and Oura Rings have become integral to our daily lives as software-based calorie counters, pedometers, heart rate sensors, stress monitors, and physical fitness trackers. These noninvasive tools have also branched into personalized medicine by providing quantitative and longitudinal data on a user's fitness or heart health [1]. As wearables become increasingly accepted as daily use devices, it is essential to explore signal sources and device paradigms that can provide new insight into our health. Electroencephalography (EEG) is one such technique that has yet to be commercialized at scale but has the potential to provide long-term information about our mindset, sleep, and focus [2]. As a noninvasive way to monitor the brain's electrical activity from the scalp, EEG has been used in the clinic to monitor chronic neurological illnesses like Epilepsy, and Dementia [3][4]. Unfortunately, recording EEG in the clinic is - to put it mildly - a pain [5]. While noninvasive, clinical EEG systems still need technicians to help users don/doff the headsets and electrodes. Additionally, these electrodes typically rely on cumbersome skin preparation and cleaning. The required clinical readout electronics are neither easy to use nor low-profile, and above all, clinical systems are not yet comfortable enough for daily ambulatory use. Fortunately, there has been growing interest in the scientific community to change that.

This report aims to introduce sensors and systems for recording EEG from inside the ear canal (Ear EEG) with user-generic, dry electrodes and wireless readouts. This manuscript will provide essential background on EEG signals, electrode modeling, dry electrode fabrication techniques, user-generic earpiece design, basic readout design principles, EEG experimental design, and an Ear EEG example application involving brain state classification.

1.2 EEG Source and Signal Breakdown

EEG arises from aggregated neuron activity

EEG signals begin with cortical neuron activity. Neurons are the fundamental unit cell of our brains and are essential for the signal transduction, propagation, processing, and translation happening around the human body. Neurons have three main parts, the soma (cell body), the axon, and the axon terminals (Figure 1.1a). As these neurons 'fire' (also known as generating action potentials), a potential difference (voltage) with respect to the neuron's extracellular environment travels from the soma to the axon terminals, effectively acting as signal amplifiers and wires for the circuits inside and around our brains (Figure 1.1b). This potential difference can also be viewed as a polarity with a specific direction (or signal vector) inside the brain (1.2a). As signals start and travel through different brain regions, they may stimulate further activity or inhibit downstream activity. Whether or not this activity is excitatory or inhibitory activity is regulated by specific brain circuits and various neurotransmitters ([6] is an excellent resource for those interested in neuroscience). Regardless of the type of signal, if the signal/stimuli are strong enough, then whole populations of neurons can 'fire' together to generate even larger voltages. The synchronized activity in large neuron populations can create large-scale changes in a brain's electric field. Often these changes are modeled as physical electric dipoles with a positive and negative charge separated by some small distance (figure 1.1b) [5]. The signal characteristics: amplitude, frequency, and vector are correlated with different actions and mental states and can be used to detect and predict actions.

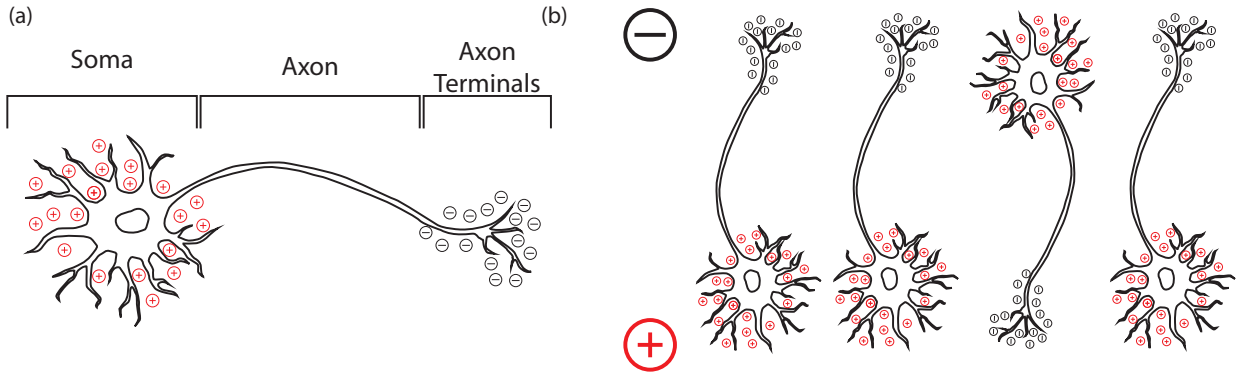

Figure 1.1: (a) Cartoon diagram of a single diagram immediately after an action potential has propagated through the neuron. (b) A group of neurons that have all fired and generated a dipole as described by the large plus and minus charges.

When inside the brain (or any more or less homogeneous medium), neural signals are freely conducted through ionic conduction. However, since the skull and tissue interfaces are not homogeneous mediums (relative to the brain), ionic conduction cannot cross these

interfaces and carry neural activity to the scalp. Ionic charges remain on either side of these tissue interfaces and form double layers. This double layer acts like a capacitor; thus, cortical signals must capacitively couple from the brain, through the skull and tissue interfaces, to electrodes placed on the scalp 1.2a. It is important to note that, unlike benchtop electrical systems, there is no explicit reference point (or ground) on the human body. Thus electrophysiological signals are differentially measured between a sense and a reference. This differential measurement is often organized to maximize a target signal's amplitude with reference to a quiet reference location on the body. For ECG measurements, reference electrodes are placed on the right leg, far away from all the sense electrodes. In clinical EEG measurements, a reference electrode is usually placed on a mastoid, a bony structure behind the ear, to maximize the distance between the scalp electrodes and the shared reference. In addition, it is also beneficial to have a third electrode to ground the body to the recording system's ground. This helps ensure the subject's body is within the recording system's input range and also reduces environmental interference (50/60 Hz from the power mains, other wireless signals permeating the air, etc.) (see chapter 5).

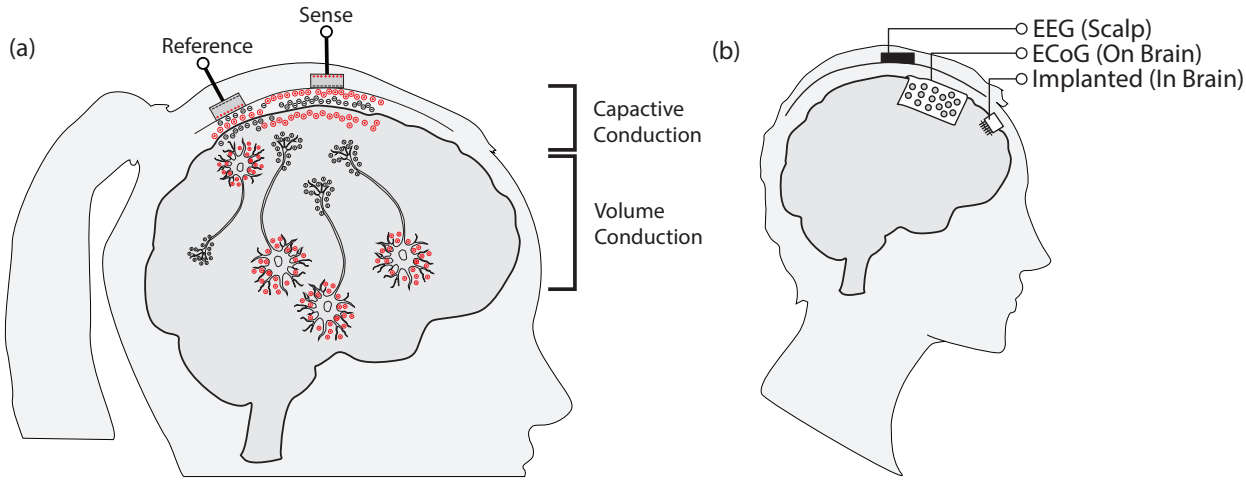

Figure 1.2: (a) Signal path via volume conduction and capacitive conduction from a group of neurons to a sense and reference electrode pair. (b) Differences between noninvasive and invasive brain monitoring techniques. The closer the electrode is to the signal generators (cortical or deep brain neurons), the fewer filters affect the signal. As a result, more invasive recording techniques can result in high amplitude signals with higher bandwidths.

EEG amplitude and bandwidth versus invasive methods

While neural signals are measurable from the scalp, scalp EEGs cannot provide the same spatial resolution, temporal resolution, or signal-to-noise ratio (SNR) as signals recorded deep inside the brain. When recording local field and action potentials (LFPs and APs) inside the

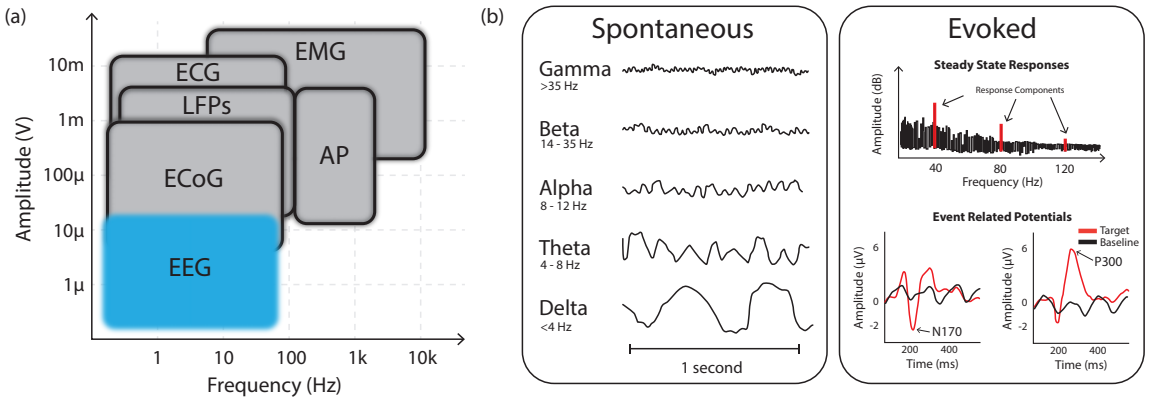

Figure 1.3: Comparison of electrophysiological signal amplitudes and frequencies. (a) Comparison of different neural signals and their average amplitudes and frequency signals. (b) Breakdown of different types of EEG signals: spontaneous and evoked potentials. Steady-state responses are most clearly observed in the frequency domain, while event-related potentials are commonly plotted in the time domain.

brain with implanted electrodes, signals can be as large as one mV and have bandwidths over one kHz. As sensors are placed further from deep brain structures, such as on the brain's surface, measurable signal amplitudes are typically less than 1mV and have lower bandwidths (100 Hz) than LFPs and APs. This trend continues, and by the time signals are recorded from the scalp (EEG), signal amplitudes are generally less than 100 μV with bandwidths less than or equal to 100 Hz (Figure 1.3a). As these neural signals propagate through tissue, the skull, and the scalp, they undergo spatial and frequency filtering from the physical averaging of neuron populations and capacitive coupling [7] [8]. Thus the further away we get from the signal source, the greater the filtering and the harder it is to record. This spatial averaging is the primary driver for the dense electrode arrays that clinical EEGs rely on. EEG SNR is so low that if an electrode is even centimeters away from the desired sensing location, the measurement may miss the target signal entirely, especially if the signal has a narrow vector through different lobes or cortical regions. Electrode placement and arrangement around the scalp matter!

To further obfuscate these signals, muscular and cardiac activity can also be recorded across the scalp and body through electromyography (EMG) and electrocardiography (ECG), respectively. EMG and ECG signals are the most prominent signals that can be noninvasively measured from the body. While figure 1.3a implies conflicting signals bandwidths where EMG signals may contaminate EEG signals, experimental design can control for different interferers.

While EMG and ECG signals may be measurable on and around the head, they are unlikely to share the same bandwidth as your scalp EEG signals. Furthermore, experiments are explicitly designed to control for undesired biosignals and take baseline measurements

without the target neural activity. In other words, if there is a large amplitude 1 Hz signal in both your control trials (that use dummy stimulus) and the target trials (with the true stimuli), it is likely ECG related (1 Hz is 60 beats per minute). Similarly, if there is increased activity above 100 Hz, regardless of the environment or stimuli, it is likely EMG related. Further examples are discussed in chapter 5.

Spontaneous and Evoked EEG Signals

There are two main classes of EEG activity: spontaneous and evoked. Spontaneous EEG signals may exist from as low as 0.05 Hz up to 100 Hz, but multiple different sub-bands are correlated with different emotional/mental states (Figure 1.3b). These bands are referred to as delta (δ), theta (θ), alpha (α), beta (β), and gamma (γ). Their associated frequencies are in figure 1.3b and table 1.1.

Table 1.1: Table breakdown of EEG frequency bands and associated mental states

EEG Band	Frequency Range	Emotional State
Delta (δ)	0.1 - 4 Hz	Sleep
Theta (θ)	4 - 8 Hz	Drowsiness
Alpha (α)	8-13 Hz	Relaxation
Beta (β)	13-35 Hz	Thinking and problem solving
Gamma (γ)	>35 Hz	Conscious perception

Evoked EEG signals not associated with these emotional states are also referred to as evoked potentials (Figure 1.3b). Unlike spontaneous EEG, evoked potentials are not easily measured without an explicit stimulus (sound, visual display, decision-making prompt, etc.). Among evoked potentials, there are two subclasses: steady-state evoked potentials and event-related potentials (ERPs). Steady-state evoked potentials can be stimulated with audio or video streams and will result in visible signals in the frequency domain (Figure 1.3b). When evoked with an audio stimulus, they are referred to as a steady-state auditory response. When using a visual stimulation, the response is referred to as a steady state visual response. ERPs are better showcased in the time domain and are generally evoked using oddball paradigms where a subject is shown various expected and unexpected prompts. The subject's EEG may display an N170 or P300 response depending on their expectation. More information on these specific experiments is in chapter 5.

Spontaneous and evoked EEG signals have been used extensively as brain-computer interface (BCI) features [9][2][10][11][12]. In summary, these signals have enabled researchers to develop BCIs to perform choice selection, cursor movement, depression monitoring, and prosthetic control [13][14][15].

1.3 Obstacles Limiting Widespread EEG Use

Current clinical and wearable setups have shortcomings that preclude them from directly translating to consumer contexts. Clinical EEG systems rely on wet electrodes (i.e., electrodes with hydrogel) to make electrode-skin contact through hair, reduce the electrode-skin impedance (ESI), and provide mechanical stability (hydrogels act as adhesives). These hydrogels dry out over time, resulting in increased interference susceptibility and signal-to-noise ratio (SNR) degradation [16]. Moreover, the electrode application process requires skin abrasion on every electrode site and results in hair loss or skin lesions [17]. Individually applying the electrodes is also time-consuming (especially on subjects with long hair) and must be performed by a trained technician [18]. In addition, clinical setups use long wires to connect electrodes to a recording module, which exacerbates interference and motion artifacts while impairing a subject's movement [19]. Advancements in clinical EEG systems (e.g., wireless recorders, mesh vests to hold wires, and caps to hold electrodes) have improved usability for inpatient care. Still, they remain bulky for day-to-day use outside the laboratory - from a practicality and visual stigma perspective (Figure 1.4a).

Recent developments in dry electrodes and recording electronics integration attempt to address these shortcomings. Dry electrodes eliminate the use of hydrogel, thus simplifying the electrode application process. Though this improves usability, the dry electrodes can result in higher impedance relative to wet electrodes [20]. Higher impedance results in not only greater interference susceptibility but also greater electrode-related noise. To lower ESI in dry electrodes, the state-of-the-art has employed microneedles, electrode fingers, conductive composites, and nanowires [21]. However, these head-worn dry electrodes are either uncomfortably large or require skin preparation. Prolonged use of these devices often results in skin irritation or lesion formation, thus limiting their use [22]. Commercial companies have incorporated electrodes and recording electronics designed to accommodate dry electrode properties into a wireless headset (Figure 1.4). While some examples, such as Figure 1.4d and e are more compact than clinical systems, these headsets still require tedious electrode preparation and/or management and cover large parts of the scalp. Furthermore, they lack the required comfort, battery life, discreetness, and motion-artifact/interference robustness for daily and public use [23]. These limitations further result in a conspicuous form factor that will likely result in social stigma.

1.4 How Ear EEG Can Enable Neural Wearables

Above all questions of robustness and usability, one limitation is abundantly clear when looking at the existing wearables in Figure 1.4. Headsets and headbands are not socially acceptable to wear out in public. The system form factor must be inconspicuous to enable widespread use of an EEG wearable. As a result, sensing EEG from inside the ear (in-ear EEG) has been proposed as a low-profile wearable modality [24]. First demonstrated in 2011, in-ear EEG promises to be a self-administered, mechanically stable, and significantly

easier EEG modality to use.

Recent work has demonstrated that many scalp EEG features can be recorded inside the ear. Both steady-state and transient evoked auditory and visual potentials (30–80 Hz) can be recorded [25]. Specific event-related potentials like N170, P100, and P300 (Figure 1.3b) have also been recorded with in-ear EEG with oddball paradigms [26]. Moreover, low-frequency neural rhythms (1–30 Hz) and electrooculography (EOG) signals (eye blinks) have been recorded [27]. Many of these features have also been used to perform sleep scoring, and seizure monitoring with Ear EEG in small-scale studies [28] [29]. While the electrode schemes of in-ear EEG recording (ear canals and outer ear) result in less spatial coverage and smaller amplitude signals relative to scalp EEG systems [30], existing demonstrations have exhibited promising results that could enable new BCIs for everyday devices using signals from the brain's temporal lobe, the cochlear nerve, and brain stem.

Mirroring clinical EEG design progression, initial demonstrations of in-ear EEG used user-specific molded earpieces equipped with wet Ag electrodes [25][30][27] (Figure 1.5a and b and Table 1.2). These customized earpieces guaranteed consistent electrode-skin contact while the wet electrode gel (and the accompanying skin preparation) reduced ESI. Though these systems were successful first steps, a widely adopted BCI cannot rely on users to perform skin preparation, and gel application due to the associated skin damage [17]. In addition to time and comfort costs, the production cost associated with custom molded earpieces is likewise prohibitive for widespread adoption. The cost of hearing aids and electrode materials can help estimate the price of a custom molded ear EEG device. For

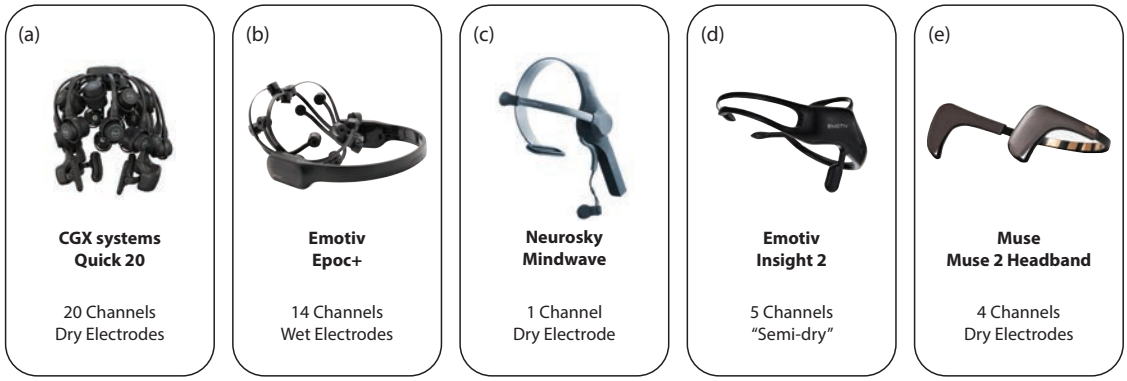

Figure 1.4: Current commercially available EEG headsets. (a) The CGX Quick 20 aims to simplify clinical EEG with a dry electrode headset. (b) The Emotiv Epoc+ is geared toward research and still requires wet electrodes. (c) The neurosky mindwave is a dry electrode headset that has been commercially sold as a toy but suffers from poor spatial resolution and low SNR. (d) The Emotiv insight's 'semi-dry' electrodes are wet electrodes with longer lifetimes. (e) The Muse 2 is a commercially sold wearable geared towards guided meditation.

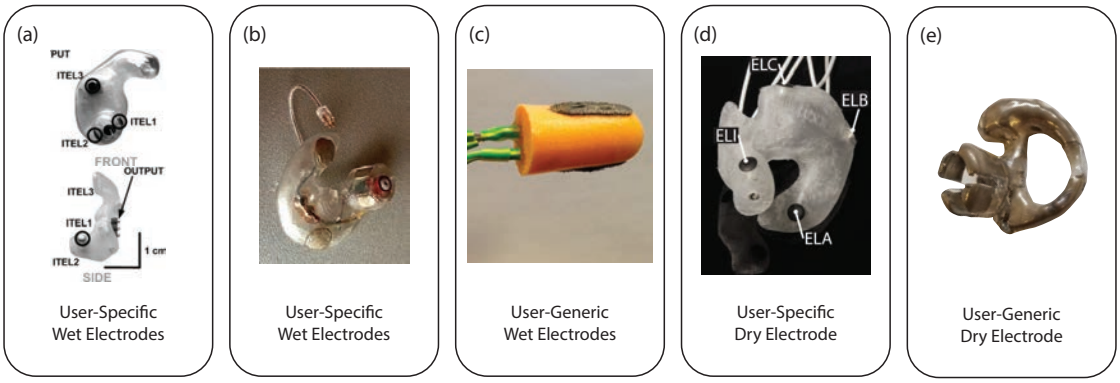

Figure 1.5: Progression of Ear EEG devices from 2011 - 2019. (a) [24], (b) [31], (c) [32], (d) [20], (e) [33]. User-generic, dry-electrode earpieces were not demonstrated across multiple users before [33].

example, the molding process alone can cost up to $200, while a pair of hearing aids can cost up to $8,000 [34]. From an electrode materials perspective, Au [35], Ag [25] [31], and AgCl [21] are commonly used electrode materials and are relatively inexpensive to acquire. However, the IrO2 used in [20] can only be purchased directly from laboratories in quantities less than or equal to five grams for thousands of dollars.

User-generic earpieces eliminate many of the fabrication costs, and as a result, later works constructed user-generic earpieces out of viscoelastic foam [32] and conductive polymers [36]. Unfortunately, they still required wet sense or reference electrodes. In addition, these earpieces relied on single-use or brittle substrates that could not be rapidly produced and used for long periods. The viscoelastic earpieces used in [32] were built off of single-use, disposable earplugs commonly used in construction sites. [36]'s conductive polymer earpieces were constructed with silver nanoparticles embedded into a silicone (PDMS) substrate. PDMS and other polymers exhibit degraded mechanical fatigue resistance and increased brittleness when mixed with large amounts of conductive particles. This is because the conductive particles physically interfere with the polymer's percolation network and limit the extent of polymerization.

There have been dry-electrode earpieces [20] built with longer-lasting acrylic plastics and Iridium Oxide (IrO2) electrodes; however, they required user-specific earpieces to maintain acceptable electrode-skin impedance and contact quality. While these dry electrodes improved usability at the cost of increased ESI, noise, and signal-quality degradation, the user-specific nature of the earpieces and the cost of IrO2 increase the barrier to entry for everyday users. Regardless, while both wet and dry electrode versions have successfully recorded spontaneous and evoked EEG signals, custom molded earpieces are not conducive to low-cost, large-scale manufacturing and user adoption. Instead, a low-cost, user-generic

solution is required.

The first user-generic, dry electrode earpieces were showcased in [33] and achieved similar performance to dry-electrode user generic designs by using large-area, compliant electrodes that applied gentle pressure on users' skin. Unlike previous user-generic earpieces, this work utilized scalable manufacturing techniques such as vacuum forming spray coating. Furthermore, it emphasized using low-cost materials by building polycarbonate and spray-coated silver earpieces. Lastly, it was also the first demonstration to use dry sense and reference electrodes. Since then user-generic, dry electrode earpieces have become more widespread [35] [37] [38].

Table 1.2: Table comparing ear EEG earpiece progression from 2011 - 2019

	[24]	[36]	[32]	[38]	Kaveh 2019
User-Generic	No	Yes	Yes	Yes	Yes
Wet/Dry Sense	Wet	Dry	Wet	Dry	Dry
Dry Electrodes	0	1	0	6	6
Wet/Dry Ref	Wet	Wet	Wet	Dry	Dry
Electrode Material	AgCl	Ag	Ag	IrO2	Ag
Electrode Area	–	–	$40mm^2$	$9.6mm^2$	$60mm^2$
Impedance at 50 Hz	–	–	–	$1.1M\Omega$	$377k\Omega$
Substrate Material	Molded Plastic*	PDMS	Foam	UC Cured Acrylic	Polycarbonate
Cost of Materials*	High	Medium	Low	High	Low

*Estimate based on material selection and fabrication procedure

These new earpieces promise to be significantly more useable and longer lasting than previous EEG wearables due to their familiar form factor, dry electrodes, and resilient material selection. No technicians are required to troubleshoot the electrode fits, and the user will not have to worry about chronic skin conditions from repeated use. Unfortunately, user-generic dry electrodes tend to have even higher ESI and record lower signal amplitudes than custom-fit earpieces. This is likely because user-generic earpieces cannot go deeper into the ear canal where there is less fatty tissue and hydrated skin closer to the signal source (the brain). As a result, it is crucial to use recording electronics uniquely equipped to deal with high impedance electrodes and low SNR signals.

1.5 Synthesizing an Ear EEG Wearable

To summarize, a discreet Ear EEG wearable requires:

1. Form factor must be discreet
2. Setup must be easy and possible without the help of technicians
3. Earpieces cannot cause any chronic pain or irritation
4. Electronics must handle high impedance electrodes and be able to measure low SNR EEG
5. System must be low power to provide long-term monitoring
6. Wearable applications have to provide meaningful utility to the user

Each subsequent chapter will focus on one of these requirements and walk through the three main blocks shown in Figure 1.6 to understand the fabrication process, electronics, and experiments required to build a user-generic dry electrode Ear EEG system. As the questions mentioned above are answered, theoretical design principles of Figure 1.6a, b, and c, and practical solutions to everyday problems that may arise through experimental use will all be discussed. These principles will include user generic earpiece design, electrode fabrication, and experimental design. After stepping through each block, methods of iterating on wearable design with the potential data gathered with an initial device will be discussed. In addition, a great emphasis will be placed on the rapid prototyping ability of different design/assembly processes. All human studies were approved by UC Berkeley's Institutional Review Board (CPHS protocol ID: 2018-09-11395).

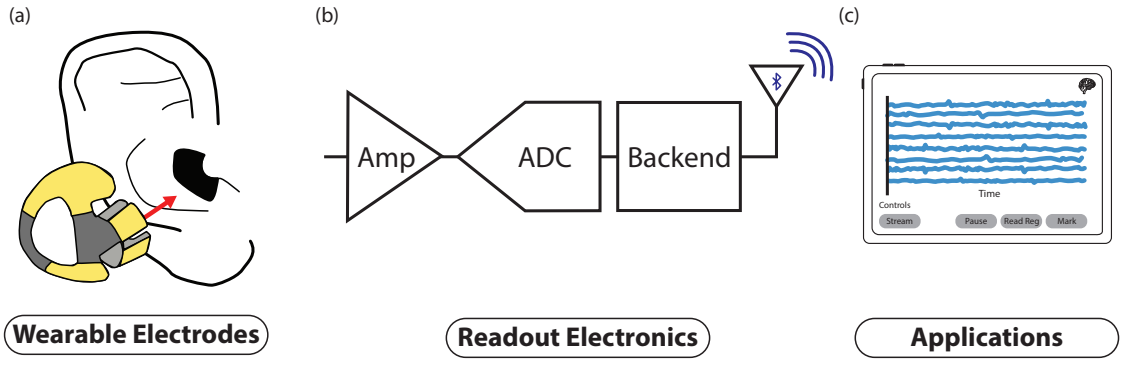

Figure 1.6: Generic Wearable System Block Diagram

Chapter 2

Electrode Modeling, Manufacture, and Use

Electrodes are electrical conductors that bridge electrical circuits with nonmetallic systems and are essential to batteries, chemical plating processes, and electrophysiological systems. This chapter will introduce electrode operating principles, a derivation of basic electrode circuit models, practical manufacturing methods, and characterization techniques to provide a broad understanding of electrophysiological electrodes. While this report focuses primarily on electrophysiological recording, this chapter will broach important considerations for stimulating electrodes to provide a more holistic understanding of the small and large signal operation of electrodes. Furthermore, electrode theory and manufacturing is an interdisciplinary topic that touches on physics, material science, and chemistry. Rather than being a complete guide to all things electrodes, this chapter focuses on the design, fabrication, and characterization of wearable electrophysiological electrodes.

2.1 Theory of Operation

Broad Electrode Background

In a biological system, electrical signals propagate via ionic charge transfer, whereas electrical systems use electrons. The electrode mediates this energy conversion via capacitive and resistive charge transfer between the physiological system (ions) and the electrical system (electrons). The choice of electrode material and coating can significantly affect charge transfer efficiency. Still, there will always be a fundamental mismatch due to the physical differences between electrons/holes and ions. For instance, while electrons and holes have mobilities around $1.35E3$ and $0.48E3$, respectively, common ions mediating charge transfer around neurons, H+, OH-, and Na+, have mobilities around 3.625E-3, 2.050E-3, and 5.193E-4, respectively. That's six orders of magnitude! This difference in mobility means the ionic current density is lower than electron current density, fundamentally affecting the way

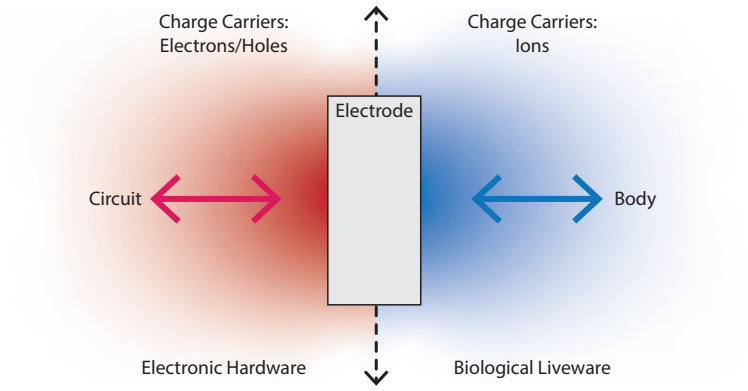

Figure 2.1: An electrophysiological electrode as a transducer at the interface between electronic hardware and biological liveware. Charge transfer is mediated by electrons in the hardware, whereas ions transfer current in the liveware.

capacitive interfaces operate across these two environments. Similarly, charge transfer inside the human body does not operate with well-defined waves through determined paths as it would through metal wires. Ions can be carried by different diffusion mechanisms in addition to voltage differences. Lastly, depending on the rate of change or signal amplitude at the electrode-electrolyte interface, the conversion can shift between small and large signal regimes dictated by different chemical processes.

The linearity of these regimes is wildly different, so it is important to understand how to stay within one or the other. Fortunately, the distinction is simple: neural recording applications deal with small amounts of voltage and current across the electrode interface. Generally, they stay within the small signal regime where capacitive charge transfer dominates. Stimulating systems must inject significantly more current and generally operate in the large signal regime where electrochemistry may occur. Resistive charge transfer begins to play a role in this large signal domain.

Electrode-Electrolyte Capacitive Double Layer

A capacitive electrode interface is formed as soon as the metal electrode is placed inside a solution/electrolyte. Certain chemical reactions will immediately occur to account for the presence of this new material. The simplest reaction will likely be the dissolution of metal ions in solution and the associated electron transfer (with a donor, D, and acceptor, A) :

$$D \rightleftharpoons A^{+} + e^{-} \tag{2.1}$$

The reaction in equation 2.1 will proceed as metal ions are formed through oxidation (a reaction where a doner loses an electron) and are consumed by reduction (where an acceptor gains an electron). Depending on the abundance of donors or acceptors, one side

will dominate until the interface reaches an equilibrium where the oxidation and reduction reactions operate at the same rate (meaning no net current across the interface). At this point, any further perturbation would be the result of external bias and the subsequent electrochemistry.

As metal and other ions adjust to this reduction-oxidation (redox) driven equilibrium, a physical space charge layer is formed (Figure 2.2a). Polar solvent molecules (water) form a hydration sheath around the metal and any free floating ion. These charged ions are attracted to the metal surface but can only get as close as the hydration sheath/layer allows, creating a minimum distance between the metal and dissolved ions. Scientists with an electrical engineering background may observe two separated planes of charge and the makings of a capacitor. The electrode's charged metal surface forms one capacitor plate in the ideal case. The hydration sheath (inner Helmholtz plane) and hydration layers create a dielectric. The attracted dissolved ions form the other charged plate (the outer Helmholtz plane) of this idealized capacitive interface. Hydrated ions will then be dispersed on the other side of this sheath (in the outer Helmholtz plane) to form the other charged plate of this idealized capacitive interface.

While it is simple, modeling this interface as an ideal double plate capacitor can yield valuable insight (Figure 2.2b and equation 2.2). For a parallel plate capacitor, the capacitance is determined by the dielectric permittivity of vacuum, ϵ_0, the relative dielectric permittivity of the medium between the parallel plates, ϵ_r, the distance between the plates, d, and the surface area of the plates, A. For both dry and wet electrodes, the dielectric can be assumed to be aqueous (electrolyte gel for wet electrodes and residual sweat for dry electrodes) with an $\epsilon_r = 78.54$ and that the Helmholtz plane distance (distance between the electrodes and outer Helmholtz plane) is about 5 angstroms. This would result in a

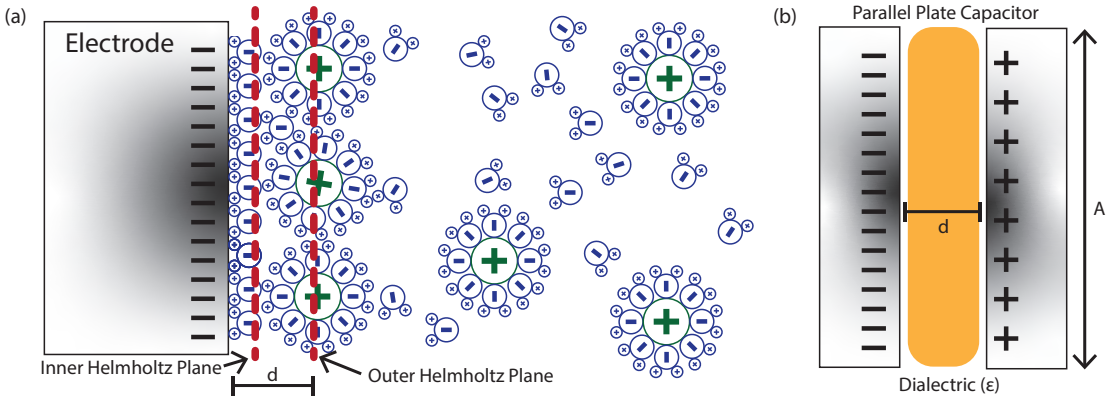

Figure 2.2: (a) Ideal electrode-electrolyte interface for an electrode in solution. A double layer formed at the interface is between the metal electrode and the outer Helmholtz plane. (b) Simple parallel plate capacitor diagram. Note that the area labeled A is the cross-sectional area not fully shown in this 2D diagram.

Helmholtz-Perrin (simplified to Helmholtz) capacitance of:

$$C = \frac{\epsilon_0 \epsilon_r A}{d} \tag{2.2}$$

$$C_H = 0.11 \frac{F}{m^2} \tag{2.3}$$

One could go even further and try to model the capacitance variation along a diffuse charge plane with differing solution concentrations using a Gouy-Chapman model [39] [40]. Still, it is essential to recognize that these models inevitably break down when designing dry electrodes to record from human skin. It becomes too complex to model the day-to-day variation in a user's sweat, skin dryness, electrode placement (skin is not uniform), or to get a fair measure of the exact contact area. Regardless, this simple model will provide intuition on how to improve electrodes and compare different electrode designs. For this purpose, the Helmholtz-Perrin model is sufficient because it gives us a physical basis for what is happening at the electrode-electrolyte/skin interface. This physical basis makes it clear that the electrode capacitance primarily depends on the electrode surface area and its average distance from the skin. Simple as that!

Resistive Charge Transfer Mechanisms

While much lower amplitude and higher frequency signals will pass through the capacitive interface, DC (direct current), some small amplitude signals, and most higher amplitude signals will prefer the resistive charge transfer path that exists in parallel to the capacitive interface detailed above. This resistive path is responsible for the initial equilibrium potential via redox reactions like equation 2.1 and nonlinearities that result from large-scale amplitudes. These cases can be driven by multiple different mechanisms that depend on the applied voltage amplitude: the reversible diffusion of redox reactants to and from the electrode, the reversible chemical reactions occurring at the surface of the electrode, the irreversible chemical reactions that occur under large voltage biases, or the dissolution of electrode metal atoms into solution.

As was the case with capacitive double layer models, there are multiple models for this resistive charge transfer path: the Butler-Volmer model and the Warburg model. The Butler-Volmer model of charge transfer models this impedance as something close to a hyperbolic sine and assumes that the linearity of resistive charge transfer is independent of an electrode materials current density and electrode area. This inherently assumes that you have a uniform electrode-electrolyte interface (large electrodes may have very different chemical reactions occurring in different regions of the electrode). While this tends to be true for electrodes in solution (wet electrodes), it is rarely true for dry electrodes. An important takeaway from this model is the hyperbolic sinusoidal shape shows that very large applied voltages are required to produce appreciable currents across the interface. This large voltage deviation from the equilibrium promotes electrochemical reactions at the electrode surface.

It may also deplete the electrode interface of reactant ions required to carry current through the electrolyte. These reactions and the subsequent ion depletion may not be immediately reversible, thus making the resistive charge transfer path highly nonlinear. For recording applications (small voltage deviations), the interface can be modeled as a linear resistor that adheres to ohm's law:

$$V = IR \tag{2.4}$$

For further information and a mathematical derivation of the Butler-Volmer model, please refer to [39][40]. This model implies that for small voltage perturbations, there exists a linear and fully reversible resistive path exists for small voltage perturbations.

The Warburg model goes a little further and models this resistive charge transfer as a constant phase element (neither a true resistor nor a true capacitor). This constant phase element (CPE), models a resistive element with a constant phase of 45° and is a better representative of the irregularities that occur across the entire electrode-electrolyte interface. CPEs are modeled by

$$Z_{CPE} = \frac{1}{(j\omega)^n Q} \tag{2.5}$$

where $0 < n \leqslant 1$. Q is a measure of the magnitude of Z_{CPE} while n fits the bilayer phase offset. Historically, Warburg impedances with CPEs were used in conjunction with other typical capacitors and resistors to make electrode circuit models. As was the case with the capacitive interface models, these resistive charge transfer models offer limited utility. Instead, they are best used to gain intuition on electrode interfaces and compare different electrode designs. The main takeaway should be that this charge transfer impedance is likely high for dry electrodes. Note that the best way to model this is to measure it in vivo and then fit a model.

Spreading Resistance

There is one last model element related to the geometric surface area of the electrode. The spreading resistance accounts for the net resistance encountered by current spreading out from an electrode into another conductor. Fortunately, it is straightforward to model and theoretically calculate.

Spread resistance, R_s, can be calculated by integrating the series resistance of shells of solution outward from the electrode:

$$R_s = \int_{x=0}^{x=\infty} dR_s \tag{2.6}$$

where x is the distance "normal" to the surface and the shape of the shell is determined by the geometry of the electrode. For a spherical source of current (coming out of a spherical electrode in solution - Figure 2.3a), R_s can be defined by:

$$R_s = \int_{x=0}^{x=\infty} \frac{\rho}{4\pi r^2} dr = \frac{\rho}{4\pi r_s} \tag{2.7}$$

Where ρ is the resistivity of the solution in $\Omega \cdot cm$, and r_s is the radius of the electrode sphere. Note that this is not material dependant! For a flat disk electrode (like an EEG electrode, Figure 2.3b), R_s is defined by:

$$R_s = \frac{\rho}{4r} \tag{2.8}$$

where r is the radius of the circle. For a rectangular electrode (Figure 2.3c), R_s is:

$$R_s = \frac{\rho ln(4L/W)}{\pi L} \tag{2.9}$$

where W and L are the width and length of the rectangle (in cm). For the case that $w = l$ (a square - Figure 2.3d), this simplifies to:

$$R_s = \rho \frac{ln(4)}{\pi L} \tag{2.10}$$

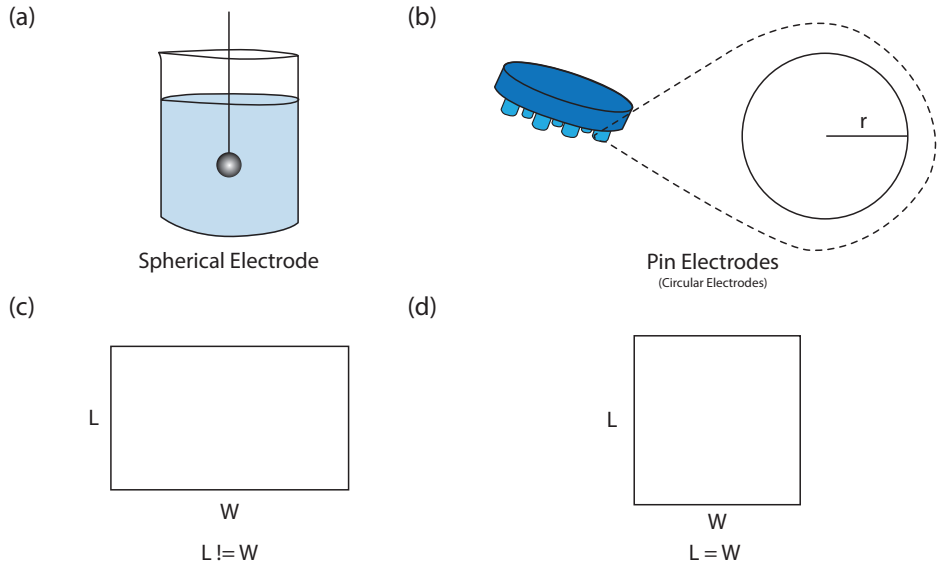

Figure 2.3: (a) Spherical electrode in solution. (b) Pin electrode with a call-out to the cross-section of an individual pin. All circular pin areas can be summed and treated as a single, large circular electrode. (c) Rectangular electrode. (d) Square electrode.

Electrode Chemistry at Large Amplitudes: Reversible vs. Irreversible Surface Chemistry

To this point, nonlinearity at the electrode interface has been repeatably referenced but not fully explained. In recording applications, electrodes primarily behave as capacitors, and no electrochemistry occurs at the surface of the electrode. In stimulating systems, this is no longer the case. When stimulating live tissue, the goal is to inject enough electric charge to depolarize a cell membrane (and promote neural activity) without killing or damaging tissue. Different materials have different inherently different reversible charge storage capacities (CSC), which dictates what charge density can be injected in a specific time frame without generating dangerous electrochemical byproducts [41]. The relationship between CSC, charge density, stimulation period, and tissue damage can be modeled by the Shannon limit relationship [42]. As a result, any reactions occurring at the electrode-electrolyte interface are ideally short-term, fully reversible, and operate below the Shannon limit for tissue damage. Reversible reactions involve ionic species that remain near the electrode surface and are not depleted during the course of the reactions. Noble metals such as platinum and iridium) are incredibly inert and have the following reversible chemical reactions for oxide formation on the electrode surface (equations 2.11 and 2.12):

$$\mathrm{Pt} + \mathrm{H_2O} \rightleftharpoons \mathrm{PtO} + 2\mathrm{H^+} + 2\mathrm{e^-} \tag{2.11}$$

$$\mathrm{Ir} + 2\mathrm{H_2O} \rightleftharpoons \mathrm{Ir(OH)_2} + 2\mathrm{H^+} + 2\mathrm{e^-} \tag{2.12}$$

When the injected voltage amplitude is too great, irreversible reactions take place and new chemical species are generated that are not bound to the electrode surface. One example reaction is the corrosion of Platinum:

$$\mathrm{Pt} + 4\mathrm{Cl^-} \rightleftharpoons [\mathrm{PtCl_4}]^{2-} + 2\mathrm{e^-} \tag{2.13}$$

The hydrolysis of water, when water is split to form hydrogen gas (H_2) and oxygen, is another irreversible reaction that can be harmful. Cyclic voltammetry can be used to empirically measure the thresholds between different charge transfer mechanisms and the voltages/charges required to start effecting new chemical species [42]. There has been significant research and work to understand how much current can pass through an electrode interface before tissue damage occurs. While incredibly interesting, it is outside the scope of this report, and further discussion will be limited to recording focused electrode design.

Impact of Different Materials

Material selection is a critical step in designing electrodes. Not only does the material have an important bearing on system performance, but it also impacts user safety and manufacturing ability. In short, there are a number of criteria that wearable designers should keep in mind when picking materials for surface use:

1. Biocompatibility from a tissue and allergy response perspective
2. Longevity and robustness of electrode-skin impedance
3. Accessibility of the materials

When it comes to user safety, certain materials will cause skin irritation and severe allergic reactions when habitually exposed to skin. Thus, any material known to cause damage with periodic exposure should be avoided. Such materials include: nickel, chromium, cobalt, mercury, and copper [43][42]. It is important to note that this list is neither exhaustive nor geared towards implanted electrodes.

Not all biocompatible metals have robust, long-term performance. Silver is one such example. While silver is commonly used in medical devices, it is primarily used for short-term recordings since the electrode oxidizes (tarnishes) over time. This oxide growth increases its impedance and hurts device performance due to the fact that most oxides are non-conductive. Most other noble metals (gold, platinum, iridium, and palladium) are excellent candidates due to their inert nature. They do not oxidize in ambient environments and are highly biocompatible [43].

Unfortunately, some of these metals are incredibly expensive for no serious benefit. Furthermore, they are difficult to buy in large enough quantities for electrode manufacturing. Palladium, while inert, conductive, and biocompatible, is one such expensive metal. Iridium and its conductive oxide form (iridium oxide) is another example of an amazing electrode material but is generally prohibitively expensive. Metals like gold and platinum tend to be the most popular candidates due to their accessibility and existing widespread use. Gold can be easily milled, evaporated, and electroplated, making it incredibly easy to work with.

As a side note, it is important to understand there is another set of safety criteria for implanted electrodes: radiographic visibility, ferrocity (how magnetic it is), and the amount of scar tissue that may be formed due to the metal/electrode size. These are important specifications to keep in mind in order to make follow-up doctor visits easier and prevent accidental harm/excessive heating in magnetic resonance imaging (MRI) machines.

2.2 Models

Generic Model

Electrical models of the electrode-electrolyte interface can be assembled using the charge transfer resistance, double layer capacitance, and spread resistance mentioned in Section 2.1. Since the interface has both capacitive and resistive charge transfer mechanisms, most models will put a capacitor, C_{DL}, and resistor, R_{CT}, in parallel with each other. Whatever current passes through these two elements must also pass through the physical electrode surface cross-section, and thus a spread resistance, R_S, is placed in series with the C_{DL} and R_{CT} pair (Figure 2.4a). C_{DL} and R_{CT} represent the double layer capacitance and charge transfer resistance, respectively.

A variation of this model comprises resistors and a CPE (in place of the capacitor - Figure 2.4b). As mentioned previously, a CPE is an equivalent electrical model for a double layer and is often used as a measure of the electrode-skin interface's non-faradaic impedance. These models are adaptations from the theoretical models typically used to represent the electrode-electrolyte impedance [44][45][46]. A standalone Warburg impedance, due to chemical reactant diffusion in solution, is not included in this model because it has been experimentally determined that the CPE-based model captures the appropriate behavior for most materials and frequencies of interest [35][47].

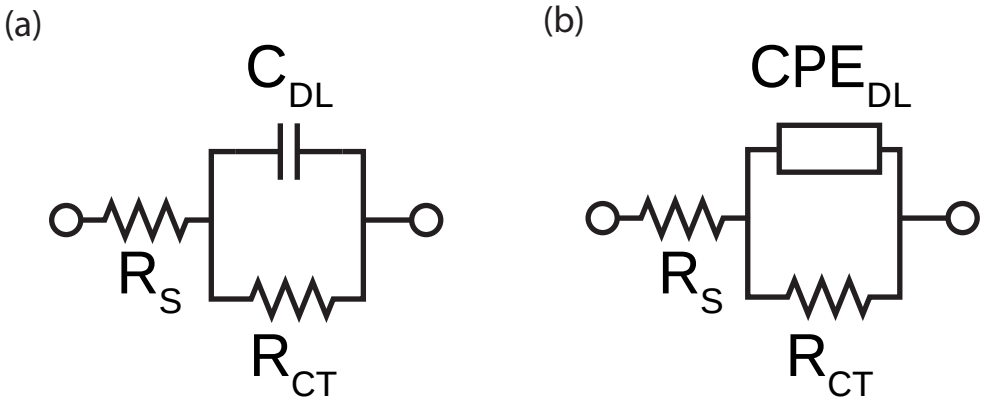

Figure 2.4: Individual electrode models. (a) Capacitive electrode model. (b) Constant phase element model.

These models represent the single electrode-electrolyte (or electrode-skin) interface, and in most cases, are sufficient for modeling and system comparison. Capacitive electrode models have the following impedance:

$$Z_{e,CAP} = R_S + \frac{R_{CT}}{1 + \jmath\omega C_{DL} R_{CT}}. \tag{2.14}$$

While CPE-based electrode models have the following transfer function:

$$Z_{e,CPE} = R_S + \frac{R_{CT}}{1 + \jmath\omega^n Q R_{CT}}. \tag{2.15}$$

These models can be analyzed further to approximate their pole/zero behavior on a bode plot. To do so, the capacitive model's poles and zeroes must be calculated and then the corresponding frequency behavior can be compared to experimental data.

First, all terms must be combined into a single fraction like the below equation:

$$Z_{e,CAP} = \frac{R_{CT} + R_S(1 + \jmath\omega C_{DL} R_{CT})}{1 + \jmath\omega C_{DL} R_{CT}}. \tag{2.16}$$

The numerator represents the zero, while the denominator represents the pole. To find the zero/pole, the numerators and denominators are set equal to 0, respectively. From here, ω can be found for both the zero condition (numerator) and the pole conditions (denominator).

$$Z_{e,CAP,zero} : R_{CT} + R_S(1 + \jmath\omega C_{DL} R_{CT}) = 0. \tag{2.17}$$

$$Z_{e,CAP,pole} : 1 + \jmath\omega C_{DL} R_{CT} = 0. \tag{2.18}$$

Solving both equations with the assumption $R_S << R_{CT}$ (as is the case for dry electrodes) yields:

$$\omega_z = \frac{1}{R_S C_{DL}} \tag{2.19}$$

$$\omega_p = \frac{1}{R_{CT} C_{DL}} \tag{2.20}$$

This implies a lower frequency pole and a higher frequency zero because our earlier assumption would also indicate that $R_S C_{DL} << R_{CT} C_{DL}$. This is confirmed by the example data shown in Figure 2.5a where there is a low-frequency pole near DC and a higher frequency zero near 2 MHz.

The practical differences between these models lie in their modeling and simulation utility. The capacitive model has a constant phase of 90° and cannot be tuned for different phases. Electrode interfaces are never this ideal and thus can have phases that range anywhere from 40° to 80°. $Z_{e,CPE}$ fits the average impedance spectrum significantly better than $Z_{e,CAP}$, since CPEs can model phase shifts less than 90° as can be seen in the example impedance

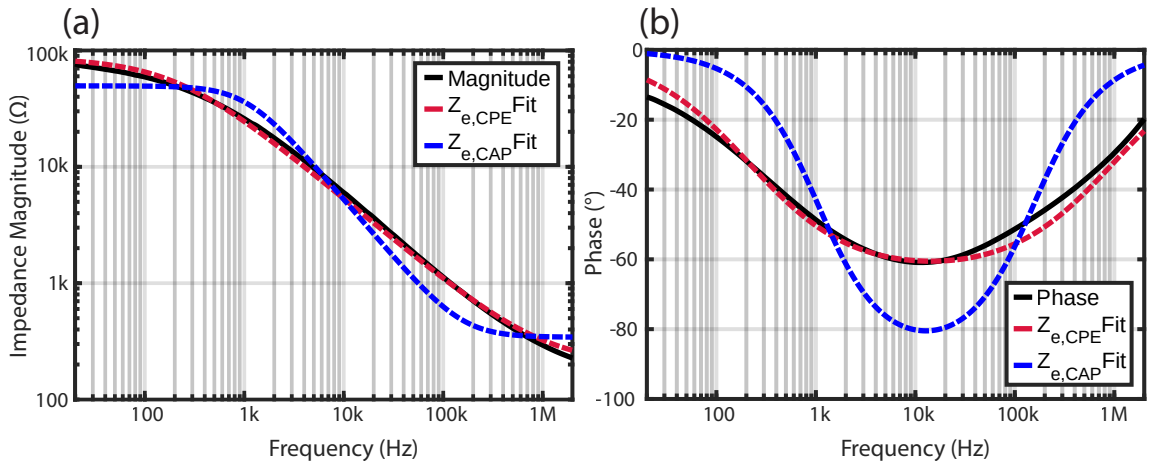

Figure 2.5: Typical values and model fits for the (a) magnitude and (b) phase of a dry electrode's impedance spectrum. CAP and CPE-based model fits represented by the blue and red hatched lines.

Table 2.1: Model Fit Parameters

C_{DL}	Q, n	R_{CT}	R_s
$3.044nF$	–	$49.27k\Omega$	343.5Ω
–	$.69n, 0.7$	$87.67k\Omega$	195.1Ω

phase (Figure 2.5b). In simulation software, where CPEs exist, the CPE model can provide more accurate results. On a bench top, CPEs do not exist as discrete circuit components. Thus, capacitor-based models tend to be more useful since they fit the electrode model well enough, and they can be used on a breadboard to model an electrode's interaction with a larger system.

In some select cases, these models can be stacked to describe every tissue interface between the body and the electrode (Figure 2.6). As will be shown later, this usually does not

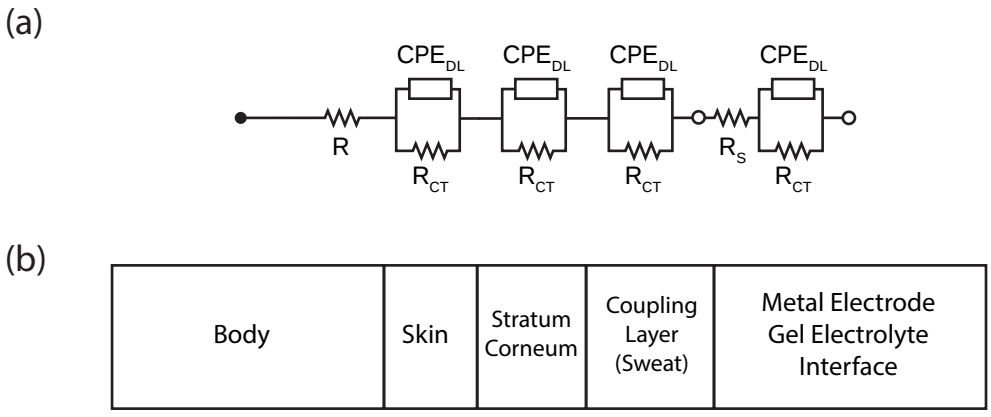

Figure 2.6: Stacked models to capture multiple tissue interfaces.

provide any extra benefit for the increased mathematical complexity. Thus the models in Figure 2.4 have remained more widely used and will remain the primary focus of this report [35][47][48].

Electrode DC Offset

Every reduction or oxidation reaction has a transfer of charge and an associated voltage. For a given redox reaction (where an oxidizer, Ox, is converted to its reduced form, Red):

$$\mathrm{Ox} + \mathrm{ze}^{-} \rightleftharpoons \mathrm{Red} \tag{2.21}$$

The voltage difference generated by this reaction is defined by the Nernst equation:

$$V_{cell} = V_{cell}^{0} - \frac{RT}{zF} ln(Q_r) \tag{2.22}$$

Where V_{cell}^{0} is the standard electrode potential that has been experimentally measured for different redox reactions, R is the universal gas constant, T is the temperature in kelvin, z is the number of electrons transferred for a given redox reaction, and F is Faraday's constant (in coulombs per mole). Lastly, Q_r is the reaction quotient for the given redox reaction and is defined by:

$$Q_r = \frac{|Red|}{|Ox|} \tag{2.23}$$

where the bars $|x|$ denote the effective concentration of generic chemical species, x. Since redox reactions occur across the faradaic charge transfer resistance at the surface of our electrode, there is a change in potential associated with the electrode interface. This is referred to as an electrode half-cell potential and is denoted by the battery symbol in Figure 2.7a (in fact, this is a very important topic in battery design as well - though that is for someone else's thesis). This half-cell potential is determined purely by the Nernst equation (eq 2.22) and is not directly related to the electrode shape, size, or surface texture. It purely depends on the chemical species undergoing oxidation or reduction at the electrode interface. This means the electrode material and the electrolyte (which may be sweat) are responsible for any voltage observed when no external bias is provided.

When a complete recording system is set up with a sense, reference, and ground electrode (Figure 1.2), the complete circuit-body interface will resemble the simplified model in Figure 2.7b. In an ideal world, each individual electrode interface would have the same half-cell potential, and all impedances would be the same. This would result in the half-cell potentials canceling each other out when probing across two electrodes. While this may be close to the case when electrodes are in an aqueous solution and full electrode-electrolyte contact is ensured, in most on-body scenarios the electrode half-cells rarely entirely cancel. Small differences between the electrode interface (the distribution of sweat, skin oil, lotion, etc.) will introduce different concentrations of redox species at the different interfaces, thus resulting in some non-zero open circuit potential, a voltage measurable when 0 current is applied across two electrodes.

This open circuit potential could be referred to as an electrode DC offset (EDO) that will always be present during physiological measurements. As long as the EDO is less than the input range of the recording instrumentation, it has no bearing on recording signal quality

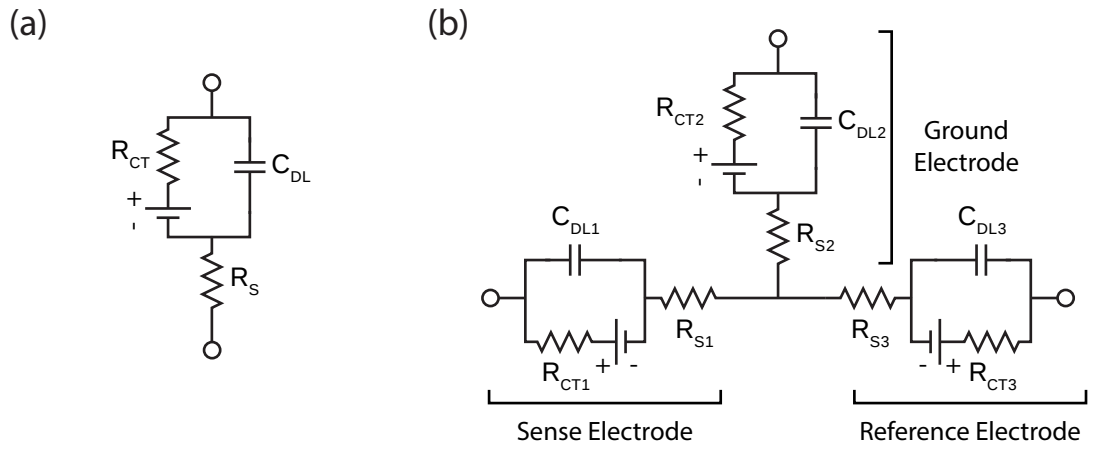

Figure 2.7: (a) Single capacitive electrode model complete with a half-cell potential by the charge transfer resistance. (b) Complete measurement model with a sense, reference, and ground electrode. Figure modified from [42].

and can be easily filtered out. Regardless, due to the variable nature of the electrode-skin interface, it is important to take many measurements on actual human skin to have confidence in any statistical measure of an EDO. These in vivo EDO measurements are crucial because it is otherwise unrealistic to fully model the EDO of a dry electrode in the presence of oil, lotion, or sweat. EDO is not constant and depends on the day-to-day nature of the interface's chemical make-up. For example, EDO may increase or decrease and settle to a new value as a subject sweats. Experimental protocols should allow for this settling time to ensure stable and repeatable measurements.

Differences between Wet and Dry Electrodes

Every electrode model and behavior described above applies to both wet and dry electrodes. Differences between wet and dry electrodes all relate to the fact that wet electrodes assume perfect and complete electrode-skin contact while dry electrodes do not.

Wet electrodes generally rely on skin preparation and always utilize a coupling electrolytic layer between the metal electrode and the skin (Figure 2.6). Typical skin preparation involves skin cleaning with alcohol (to dissolve and clean any oil or grime) and skin abrasion (to remove any dead skin that may interfere with ESI or electrode adhesion). While still considered non-invasive, these procedures can result in skin irritation, hair loss, and skin lesion formation [17]. Furthermore, a trained technician must perform these preparation steps [18] [19]. From a modeling perspective, the electrolytic gel and freshly inflamed skin ensure an ample supply of ions for charge transfer. This results in a lower charge transfer resistance for wet electrodes compared to those of dry electrodes.

Dry electrodes rarely/if ever, make perfect contact with the skin. This is because the skin is not a feature-less, flat material. Human skin has a rough surface with pores, waves, and cracks. This inherent non-linearity and the corresponding reduction in electrode-skin contact can result in an increase in ESI (for wet and dry electrodes) of as little as 30% to entire orders of magnitude [20][49]. In addition to this rough surface, uncleaned skin will have non-conductive dirt and oil on its surface. These non-idealities will result in inconsistent contact areas between the sense and reference electrodes as well as different chemical species at each interface, making each electrode's half-cell potential and impedance highly dependent on random variables like time of day, lotion use, whether or not the user recently worked out, and more. This randomness results in highly variable EDOs and ESIs. To describe this difference between theoretical and experimental electrode performance, researchers will distinguish between an electrode's geometric area and effective area. An electrode's effective area accounts for surface roughness and the presence of hydrogel. The effective surface area also directly impacts and electrode's charge transfer resistance and double layer capacitance.

Ultimately, dry electrodes have a wider spread of model parameters. Depending on the specific user and time of day, the R_{CT} and EDO may vary by the same magnitude of their average values [35]. Wet electrodes have overall lower ESI values and more consistent/controlled EDOs.

Practical Design Considerations

How can designers combat the inherent variability in dry electrode behavior? Fortunately, there are a couple of design considerations: effective contact/surface area, surface chemistry (material), and electrode design (material/design).

The electrode models shed light on some important correlations that designers can take advantage of:

- An electrode's R_{CT} and C_{DL} are directly proportional to the electrode area. R_s is directly proportional to the geometric area.
- EDO is primarily dependent on material choice and other chemical species at the interface
- Variability in electrode behavior is driven by variations in electrode use and wear

The electrode area can be increased by making the electrode physically bigger, but that may not always be viable from a recording site perspective. However, increasing an electrode's surface roughness can achieve a similar effect without increasing the electrode's geometric area. Rather, increasing surface roughness increases the electrode's effective surface area because the electrode will comply better with rougher human skin and be able to take better advantage of any residual sweat, thus reducing overall ESI.

While surface chemistry does not depend on electrode shape or size, a larger electrode will average any non-idealities to achieve similar behavior as another large electrode on the

same subject. This indirectly helps equalize the electrode half-cell potentials and minimize EDO. In addition, the individual electrodes should be made with the same shape, size, manufacturing technique, surface treatment, and material. Furthermore, inert materials such as gold or platinum should be used to reduce the likelihood of forming new oxides that dissolve in the solution.

The electrode design should also account for user habits. In the case of an earpiece, it is important to generate a design that can only fit in an ear one way. This ensures consistent electrode arrangements and user behavior, making distinct recordings similar across different days.

2.3 Electrode Manufacturing Overview

Common Surface Electrode Styles

As previously stated, clinical EEG, EMG, and ECG devices (generically referred to as ExG) rely on a low-impedance interface between conductors and skin, achieved by extensive skin preparation and application of electrodes with wet hydrogel. These surface electrodes are often referred to as cup electrodes and look like Figure 2.8a. They have a conductive metal or polymer cup that can be filled with electrolytic gel and placed on the skin. With the proper preparation and placement, these electrodes can sense through hair and on any external surface. However as mentioned earlier, despite their low impedance and high-SNR interfaces, these electrodes are neither user-friendly nor easy to don/doff for users.

Some have attempted to expand ExG to everyday users by incorporating general-purpose semi-dry and dry electrodes into wearables. Semi-dry electrodes use less hydrogel that is either pre-applied or stored in an on-electrode reservoir [53] [54]. These semi-dry electrodes can achieve similar ESIs but still require electrolytes and can suffer from control issues such as over-release. If over-release were to occur, these semi-dry electrodes also risk bridging, i.e. gel-induced short-circuits between electrodes. In addition, these reservoirs make semi-dry electrodes bulkier than wet electrodes [55]. Thus semi-dry electrodes provide small and relatively less-significant benefits than wet electrodes. Fully dry electrodes further increase usability and patient comfort but generally result in higher ESI ($> 1M\Omega s$ at $< 250Hz$) relative to wet-electrodes (10-100s $k\Omega s$ at $< 250Hz$) [54]. Microneedles, pin electrodes, conductive composites, and conformal electrodes have been used to lower ESI and improve the mechanical stability of dry electrodes [21] [36] [56] [48], however, these electrode designs

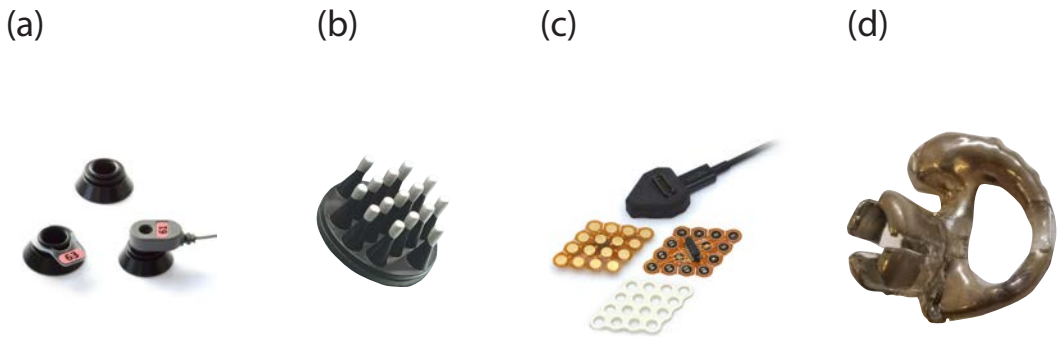

Figure 2.8: Different Electrode Form Factors. (a) Cup Electrodes [50]. (b) Pin electrodes [51]. (c) Flexible array electrodes [52]. (d) Vacuum formed electrodes [35]

introduce new use and fabrication trade-offs. Microneedles, which pierce the top layer of skin, achieve lower ESIs and enable higher SNR recordings. However, prolonged use of these electrodes may result in lesion formation and introduce the risk of infection.

Non-contact and conductive composites such as silicone carbon black and silver-glass silicone show the opposite trade-off, achieving greater comfort but suffering from higher ESIs relative to other dry electrodes [21] [36] [48]. There also exist conductive composites that have metalized tips (Figure 2.8b). While these metalized tips improve ESI, demonstrated conductive composites still have shorter lifetimes due to how conductive particles harm the elastic substrate's percolation network. These composites tend to be more brittle than their non-conductive counterparts and fatigue significantly quicker. Other electrode arrays have used flexible planar structures, machined metals, or metal printed devices to increase electrode compliance, comfort, and possible sensing location [35] [57] [58] [59]. Printed flexible electrode arrays enable high density electrode placement, high resolution 2D designs, high-volume fabrication without vacuum, and comfortable electrode compliance along a single axis such as around arms [59] [60] (Figure 2.8c and d). Recent work has also focused on making chemically inert arrays out of laser-sintered gold. These 2D arrays are potentially more robust than previous demonstrations of silver and graphene-based flexible arrays but still generally require hydrogels to record through hairy surfaces [61] [62]. 3D printing and micromachining of insulating and conductive materials have been used in combination with vacuum deposition processes such as sputtering, or evaporation [63] [64] to achieve low impedance scalp electrodes. Ultimately, existing techniques emphasize that no existing single dry-electrode design can target all biopotential signals. Furthermore, existing fabrication techniques for comfortable, anatomically customized electrodes are neither low-cost nor scalable or are limited to specific electrode shapes. As a result, highly adaptable prototyping methods with minimal tooling requirements are fundamental when beginning to design wearables.

Rapid Prototyping Methods vs. Scalable Production Methods

While very flexible prototyping techniques should be prioritized when building a new wearable, keeping scalability as a peripheral concern can help ease downstream problems. For example, micro-machining metal electrode tips out of Iridium for conductive composite electrodes (something manifold) may lead to an incredibly useable and comfortable electrode, but the amount of tooling, different steps, interconnects, and boutique materials will make these electrodes time-consuming and financially expensive to make and test. Prototype devices made from vacuum molding and 3D printing with a minimal number of materials that can be directly hot soldered, on the other hand, can be easily scaled to high-volume, low-cost production using commercial injection molding, flexible printed circuit board (PCB) production, and manufacturing techniques (assuming the device design allows for that).

A lot of the muscle memory required to know which techniques can be used for what designs must come from experience, but there are a couple of guiding principles that will always be useful. Primarily, the fewer the pieces/materials, the better. Second, machining

and vacuum forming methods can not create all shapes and sizes and cannot be relied on to make complicated forest or overhang structures. Thirdly, interconnects and other electrical interfaces between conductors can be one of the highest risk points of failure for any manifold device.

In the subsequent sections, three different manufacturing methods will be discussed alongside their strengths and weaknesses. Due to the emphasis placed on rapid iteration, these fabrication techniques are geared primarily towards rapid prototyping that can be easily scaled to large volume production with minor augmentations. The devices and processes described are the product of many collaborations with Justin Doong, Karthik Gopalan, Natalie Tetrault, Carolyn Schwendeman, Leslie Pu, Julian Maravilla, Professor Michel Maharbiz, Professor Miki Lustig, and Professor Ana Arias [33] [35] [65].

2.4 Flex PCB Electrodes on Compliant Surfaces

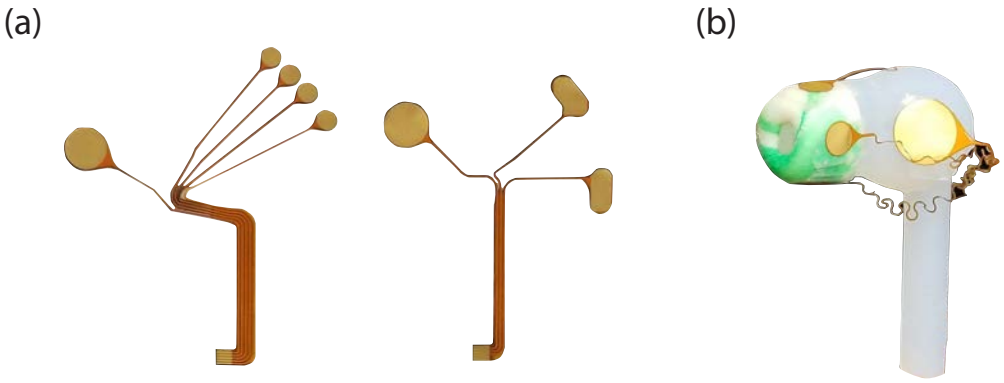

Figure 2.9: (a) Photograph examples of flex PCB electrode arrays. (b) Flex PCB arrays applied to the surface of a 3D printed Apple AirPods shell.

Flexible PCBs allow users to take existing devices (commercial headphones, watches, glasses, etc.) and add electrodes to them. Commonly used as an alternative to the more common fiberglass PCB, flex PCBs can be designed to fold inside devices or around curved surfaces. They also offer the benefit of being directly solderable and can be designed to plug into the existing flexible printed circuits (FPC) and flat ribbon cable connectors. In Figure 2.9a and b, example flex-PCB electrode arrays showcase how different gold electrodes can be arranged in arbitrary designs and while maintaining elbow FPC connectors. Figure 2.9c showcases how serpentine traces can be used to create light springs like a spring lanyard.

These serpentine traces enable not just physical flexibility to bend around corners, but also freedom in electrode placement and orientation.

Flex PCBs are generally gold plated by default providing an inert electrode interface. Flex PCBs are also made with Kapton (also known as polyimide) as the substrate since it is a biocompatible polymer commonly used to encapsulate implanted devices [66]. As a result, a fantastic feature of flex PCBs is that, by default, they provide a biocompatible and long-lasting electrode array that can be easily manufactured at scale if need be.

However, at larger electrode sizes ($> 1.5\ cm$ in diameter), the electrode is significantly less flexible. This rigidity makes it hard to use flex PCBs to place electrodes on tightly curved surfaces. Furthermore, many adhesives do not adhere to Kapton (by design), and thus integration specialists will need to resort to more clever integration solutions. While flex PCBs can be rapidly commercially produced, it is hard to produce them in-house. As a result, there may be longer lead times that may reduce a designer's iteration frequency. To get around this, designers should design multiple candidate designs and send them all out for manufacturing at once (design in bulk).

2.5 Vacuum Formed and Spray Coated Electrodes

Thermoforming and spray coating is a low-cost and repeatable manufacturing process commonly used in producing plastic shells and casings (a sample process diagram is shown in Fig. 2.10). This thermoforming-based process also enables the rapid molding of multiple devices in a single step with minimal tooling, making it an ideal rapid prototyping method.

Device construction begins by 3D printing a heat-resistant master mold with an ultraviolet cured polymer resin (Formlabs, Tough resin FTOTL4) (Fig. 2.10a). Next, a 30 mil sheet of polycarbonate, selected for its flexibility and low reactivity with isopropyl alcohol (IPA), is heated past its glass transition temperature and thermoformed around the master mold using a vacuum form chamber (450DT, Formech) (Fig. 2.10b). The master mold is then removed from the polycarbonate (Fig. 2.10c), leaving a hollow shell that is compliant and provides space for wire routing. After thermoforming, the earpiece is sanded (Fig. 2.10d) to increase the effective surface area of the substrate, cleaned with compressed air, and then placed in an IPA bath. A laser-cut polyimide (PI) mask is then applied to the device (Fig. 2.10e), which is then plasma treated with nitrogen to increase the substrate's surface energy and improve surface adhesion. The treated, masked shell is spray-coated with Ag nanoparticles (Novacentrix PSPI-1000) to deposit 15 μm thick electrodes (Fig. 2.10f). Finally, the PI mask is removed (Fig. 2.10g), and wires are cold-soldered to each device with a heat-cured silver epoxy (Epo-tek, H20E-D). Each silver epoxy bump is passivated with an ultraviolet cured epoxy. The final earpieces can be easily cleaned with IPA and reused without loss of structural integrity. This process can be easily repeated through the reuse of a heat-resistant master mold and quickly adapted to different designs, thermoplastics, electrode shapes, and materials. Acrylic, for its rigidity, may be used for the substrate, while other metals (e.g., Au for longevity) can be electroplated on the Ag base layer.

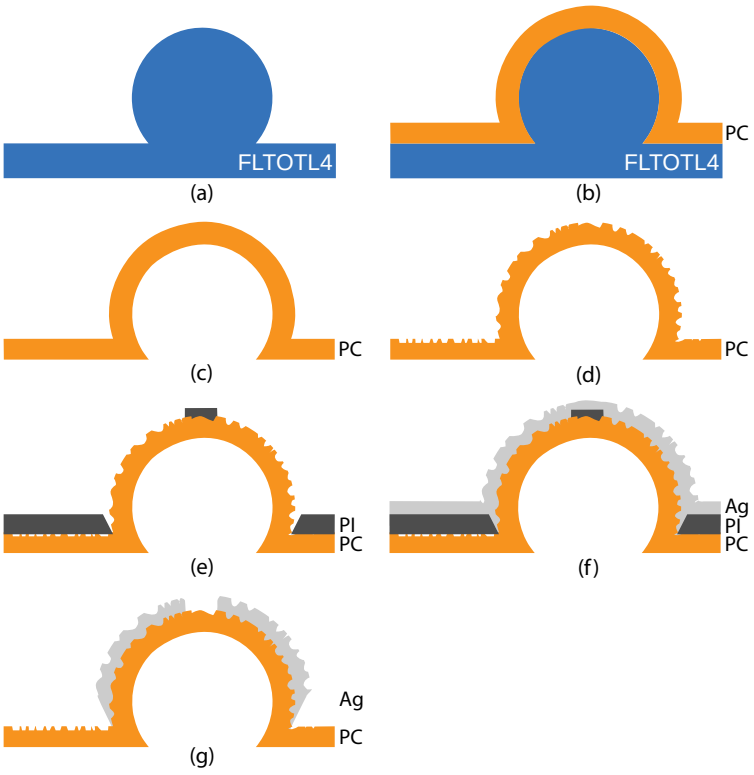

Figure 2.10: Manufacturing process diagram cross-section. (a) 3D-printed master mold (b) Polycarbonate (PC) thermoformed PC around mold (c) Hollow PC liberated from the master (d) Sanded PC after alcohol bath (e) Nitrogen plasma treated earpiece with polyimide (PI) mask (f) Ag spray-coated PC (g) PC with independent electrodes.

2.6 3D Printed and Electroless Plated Electrodes

3D printing can offer even more flexibility for electrode designs. Unlike thermoforming or thin films, designers can customize 3D printed electrodes to any anatomical surface regardless of concavity, convexity, or overhangs. 3D printing with multiple materials also opens the door to manifold devices made with both soft and rigid bodies.

The electrode structures are first printed using a stereolithography 3D printer (Formlabs Form 2 Printer) with a standard, clear methacrylate photopolymer (Figure 2.11a). Stereolithographic (SLA) printers create 3D structures by precisely laser-curing photosensitive polymer resins in a layer-by-layer fashion. This method of 3D printing results in much finer resolutions and wider choices of materials than fused deposition modeling (FDM) printing, which comprises melted plastic filaments. After printing, samples are post-processed with a 20-minute IPA bath to rinse uncured resin, followed by an hour-long UV curing process to cure the surface thoroughly.

The structures should subsequently be sandblasted with 100 grit white fused aluminum oxide blasting media (Industrial Supply, Twin Falls, ID) to increase their surface area (Figure 2.11b). Sandblasting promotes better film adhesion and lower skin-electrode impedance. The samples should then be cleaned via a sonic bath of DI water with Alconox cleaning solution for approximately 10 minutes before rinsing with DI water. The surface energy of the printed structures is then modified in a bath of 1% benzalkonium chloride (Sigma Aldrich 12060-100G) surfactant solution for 10 minutes. These steps ensure the plating surface is clean and promote the adhesion of the catalyst.

The metal plating process is a result of subsequent plating and cleaning steps. Prior to each plating step, the samples should be rinsed and dried thoroughly. First, 3D printed electrodes should be submerged in a beaker of a palladium-tin catalyst for 10 minutes, followed by the copper plating solution for a minimum of 4 hours, which provides a thick base layer of metal for gold plating as shown in Figure 2.11c. After copper plating, electrodes should be soaked in surfactant and catalyst solutions and then moved to a gold plating solution for approximately 15 minutes (Sigma Aldrich 901670-250ML) (Figure 2.11d). The second gold plating step is achieved by placing the samples in the same surfactant, catalyst, and gold plating solutions to form the top layer of the dry electrodes (Figure 2.11e). After all plating steps, tinned copper wires can be directly soldered to the electrode surface to facilitate electrode integration with recording systems. This fabrication process ensures two 0.25 µm thick layers of gold and at least 0.5 µm of copper [67] [68]. The thickness of these two layers ensures that the plated surface acts effectively as a short circuit and that any electrode impedance would be dominated by the interface and skin itself. Detailed preparation instructions and processing for catalyst and copper plating solutions can be found in Table 2.2.

Both the catalyst and copper plating solutions can be made in-house. The catalyst solution can be prepared between $60 - 70^{\circ}$ C and stirred for approximately one hour after all components are added. This solution can be prepared in full and lasts several weeks before the salts precipitate and are no longer usable. The two main components for an electroless

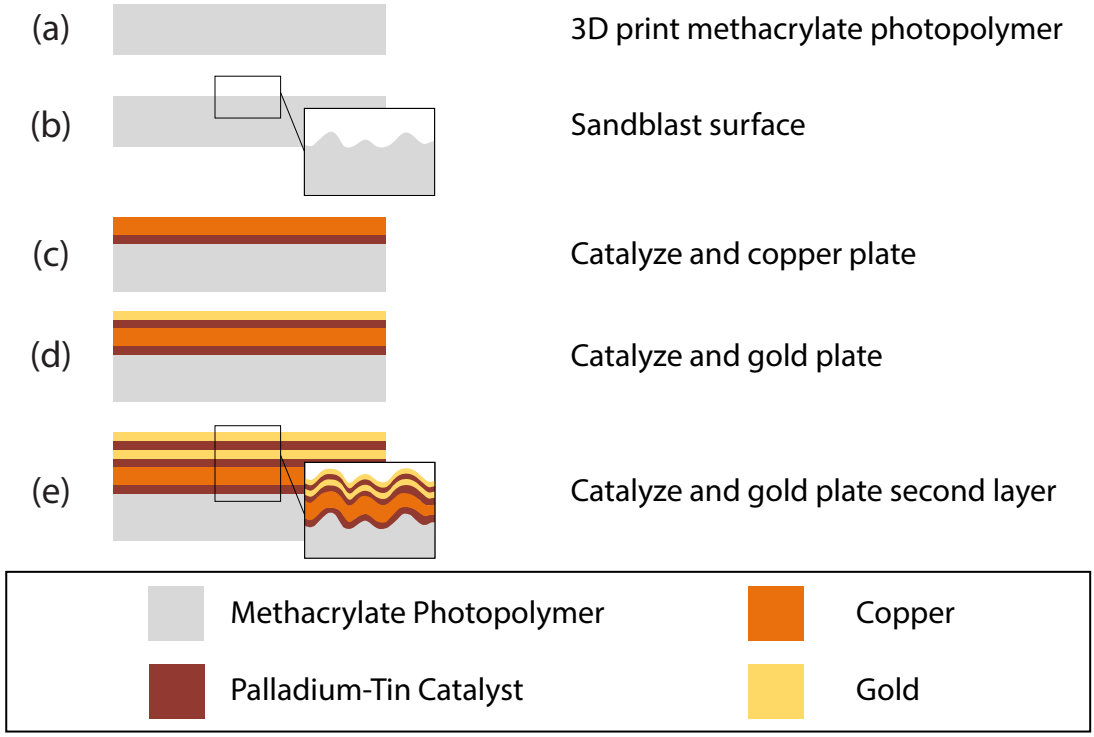

Figure 2.11: 3D printing and electroless plating process. (a) An electrode design is 3D printed with an SLA printer. (b) Samples are sandblasted and then cleaned. (c) Electrodes are metalized with copper via exposure to surfactant, catalyst, and copper plating solutions in sequence. (d) The first layer of gold is formed by submersion in surfactant, catalyst, and gold plating solution. (e) The previous step is repeated once more to improve the device's longevity and performance.

plating solution are a metallic salt and a reducing agent, in this case, copper(II) sulfate and formaldehyde, respectively. At a sufficiently high pH (the solution is adjusted to 12.8 by adding NaOH), formaldehyde reacts with hydroxide ions in the solution to reduce copper ions in the salt (equation 2.24).

$$Cu^{2+} + 2HCHO + 4OH^{-} \rightleftharpoons Cu + H_2 + 2H_2O + 2HCO_2^{-} \tag{2.24}$$

EDTA is added to the solution as a complexing agent (as copper salts are insoluble at pH greater than 4), and ferrocyanide, which acts to stabilize the solution over time (the solution is stored without the addition of formaldehyde, which is then added to the appropriate quantity being used for plating). Layers are typically built at a rate of about $1\mu m/h$ at ambient temperature, so leaving the samples for several hours or overnight in a covered plating solution provides sufficient coverage. A lightly bubbling nitrogen line was left in the

solution to provide light agitation, promote even coverage by displacing the hydrogen gas product, and limit oxidization of the copper during the plating process.

Lastly, the gold layer was applied by heating the gold plating solution to about $90°C$ and submerging the samples for about 15 minutes. This process is self-limiting, as gold layers adhere to the copper only and not upon itself. Solution recipes can be found in Table 2.2.

Table 2.2: Electroless Plating Recipes

Solution	Recipe	Notes
Catalyst	$1L$ deionized water	
	$60mL$ HCl	
	$0.25g$ $PdCl_2$	
	$12g$ $SnCl_2$	after $PdCl_2$ dissolves
Electroless Copper	$1L$ deionized water	
	$18g$ $CuSO_4 \cdot 5H_2O$	
	$48g$ $EDTA$	
	$57.2mg$ $K_4Fe(CN)_63H_2O$	
	$1mL$ HCl	
	$NaOH$	As needed to achieve pH of 12.8
	$22.5mL$ Formaldehyde	Add at plating time
Electroless Gold	Bright electroless gold plating solution (Sigma-Aldrich 901670)	

2.7 Characterization Techniques

After device construction, it is imperative to perform mechanical and electrical characterization. Understanding the pertinent mechanical and electrical specifications of resultant devices provides insight into electrode performance, lifetime, and failure modes. The goal of this section is to provide a rundown of different types of characterizations and how to perform them. Not all experiments may be pertinent to all types of electrodes, nor is this an exhaustive list. Ideally, it provides an appropriate starting point for most engineers.

Mechanical and Physical Characterization

Microscopy

Microscopy is the technical term for using a microscope to view objects or features that may not be easily seen with the naked human eye. This close-up visual inspection of newly prepared devices is paramount and can help immediately diagnose many potential electrical faults. Figure 2.12 shows a microscope image of inkjet-printed silver wires on a flexible substrate. Cursory inspection can show non-uniformities that can hurt wire conductivity (the spot in the right-most trace). Further visual inspection would also show the massive open circuit across all wires (highlighted by the red box).

Figure 2.12: Sample microscope image of inkjet printed electrode traces. Close visual inspection can illuminate fabrication errors such as the open circuits in the outlined region. Image taken with 10x magnification.

While microscopy may not help diagnose all potential faults, it serves as one of the easiest and quickest litmus checks for newly minted devices.

Surface Roughness

Measuring surface roughness (Figure 2.13 is such a surface) can be done with profilometry. There are two common types of profilometers: contact and optical profilometers. Contact

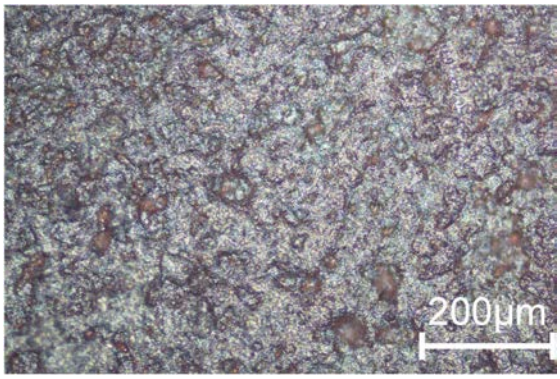

Figure 2.13: Sample microscope image of rough electrode surface. Microscopy can be used in conjunction with profilometry to measure the differences in surface topology. Image taken with 40x magnification.

profilometers drag a needle across a surface and measure a 2d depth profile, while optical profilometers will measure reflections of different lights to determine different surface roughnesses. Optical methods are generally more susceptible to different debris (dirt, water, and oil) but can measure large areas significantly more quickly than contact-based methods that have to slowly and physically drag a stylus across a surface.

There are multiple different measures of roughness, primarily R_a, R_q, and R_z. These parameters are defined below:

$$R_a = \frac{1}{l_r}\int_{x=0}^{x=l_r} |z(x)|dx \tag{2.25}$$

$$R_q = \sqrt{\frac{1}{l_r}\int_{x=0}^{x=l_r} z(x)^2 dx} \tag{2.26}$$

$$R_{z_i} = R_{p_i} + R_{v_i}; R_z = \sum_{i=1}^{n} R_{z_i} \tag{2.27}$$

R_a is the arithmetic average of profile height deviations from the mean, where l_r is the length of the measurement line and $|z(x)|$ is the magnitude of the deviation from the mean. R_q is the root mean square average of the profile height deviations from the mean (all variables are the same as in R_a). R_z is the average maximum peak to valley height of the profile within a single sampling length. R_p is the maximum peak height, while R_v is the maximum valley height. Depending on the specific device, different measures may provide additional insights. In many cases, R_a may be enough to compare different process techniques from a surface roughness perspective.

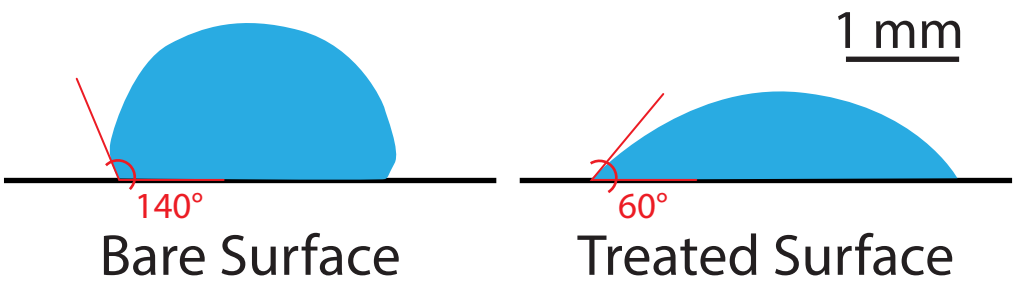

Figure 2.14: Diagram showing the difference between low and high surface energy samples. A bare surface with low surface energy will wet poorly. This results in a large contact angle. A surface with high surface energy will result in strong wetting, highlighted by the smaller contact angle.

Surface Energy

Surface energy measures how a substrate's surface disrupts intermolecular bonds on the surface of a liquid sitting on the substrate. A droplet of water will exhibit high surface tension on a smooth surface and create a tight ball (bare surface case in Figure 2.14). In this case, the smooth surface is referred to as having low surface energy with poor wetting. Wetting is a measure of how spread out the water is on the test surface. Surfaces with high surface energy will promote wetting (i.e., will disrupt water's surface tension), and a droplet of water will spread across the substrate surface (as in the treated surface in Figure 2.14). A quick and dirty comparison of surface energy can be performed by observing the change in a droplet of water's contact angle on a surface before and after treatment. Figure 2.14 displays a cartoon of such a comparison. High contact angles indicate poor wetting and low surface energy, while low contact angles are signs of good wetting and higher surface energies. In summary, rough surfaces exhibit higher surface energy, so this contact angle measurement can describe one surface's roughness compared to another.

Spring Constant

Electrode compliance determines user comfort and contact quality. A stiffer electrode may be less comfortable, but will apply greater pressure to the skin and lower ESI. As a result, an electrode's spring constant can be used as quantifiable metric to balance user-comfort and ESI. Measuring a device's spring constant involves pressing or pulling the device a certain distance with a known weight (or pressure). This force can be driven by mass and gravity (Figure 2.15) or can be applied externally by pressing on the spring (assuming the device orientation allows for it). If the device can be pressed downwards, a scale placed underneath

the device can measure the 'weight'. From there, the distance that the spring has moved must be measured. With the distance traveled, x, and applied force, F, measured, the spring constant, k, can be calculated using the relationship in equation 2.28.

$$F = kx \tag{2.28}$$

It is important to note that sometimes springs (like plastic levers) may not have a uniform spring constant. As a result, it is essential to report a spring constant at a given strain (strain being a measure of a material's deformation due to some stress). Strain is the ratio of the change in length to the original length. If you push a lever halfway from its resting position to its max position, that would be 50% strain.

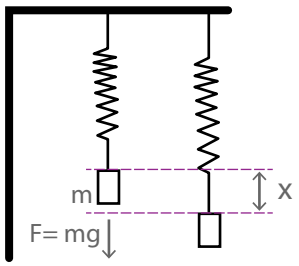

Figure 2.15: Spring force diagram. Knowing the magnitude of the force applied to a system can be used to calculate the spring's spring constant.

Electrical Characterization

Material Resistivity

While microscopy can provide information on the uniformity of your device's metal layers, there is no substitute for electrical characterization. Samples can appear continuous while still highly resistive. Thus to be confident in an electrode's overall performance, some base measures of conductivity are required.

For a three-dimensional conductor, its resistance can be modeled by:

$$R = \rho \frac{L}{Wt} \tag{2.29}$$

Where ρ is material resistivity (the quantity we're after), W is the cross-sectional width, L is the length of the conductor, and t is the sheet thickness. To achieve a dimension agnostic term, this resistance equation is commonly written as:

$$R = R_s \frac{L}{W} \tag{2.30}$$

Where R_s is known as the conductor's sheet resistance and defined by:

$$R_s = \frac{\rho}{t} \tag{2.31}$$

in $\Omega/\square$s where $\square$ is L/W. R_s does not rely on the width or length of a sample and is normalized for thickness. This makes R_s an easy-to-measure quantity commonly used to describe different conductors.

Measuring R_s involves making a 4-point impedance measurement (Figure 2.16a), where an impedance measurement system forces current through two probes and a device-under-test (DUT) and two additional probes measure the voltage across the DUT (4 probes total). Sometimes referred to as a 4 terminal (4T) measurement, this method measures just the

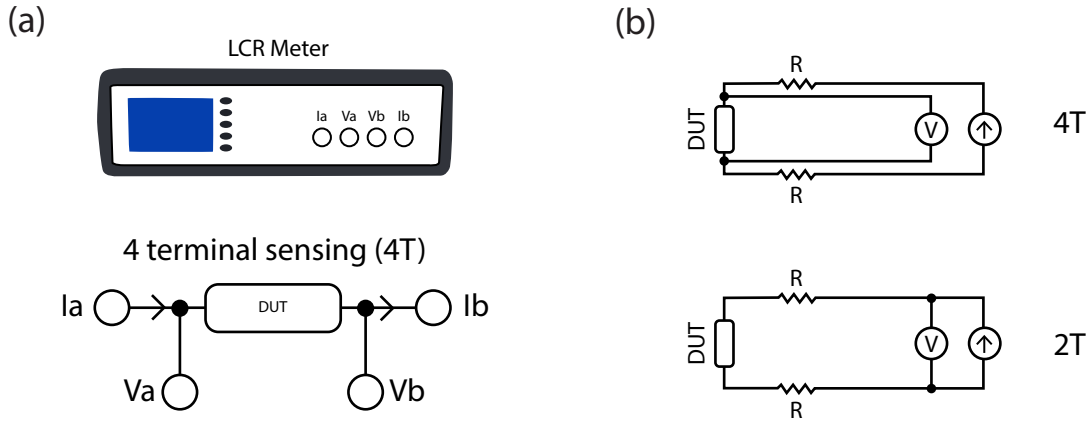

Figure 2.16: Material Resistivity and Impedance spectroscopy Test Setups. (a) LCR meters can be used to perform 4T and 2T measurements. (b) 4T measurements eliminate the electrode contact impedance while 2T measurements do not.

impedance of the DUT without measuring the contact impedance (Figure 2.16b), as would be the case in a 2T measurement. This is because the output impedance of the current test source (R in Figure 2.16b) is not in line with the voltage meter probes.

Using Ohm's law

$$V = IR_{DUT}, \tag{2.32}$$

the DUT's total resistance can be determined. From here, Van Der Pauw's method can be used to calculate the sheet resistance (assuming the four terminals are in a straight line):

$$R_s = \frac{\pi R}{ln(2)}. \tag{2.33}$$

This approximation works, especially if the goal is to measure sheet resistance changes by augmenting fabrication steps or exposure to different environments.

Electrode Skin Interface Impedance

Once a material's characteristics have been measured and designers are confident in its conductivity, the next step is to measure the electrode-skin interface impedance (ESI). Specifically, ESI is the impedance through two electrodes in contact with a human body (like in Figure 3.4a and b). The best way to encapsulate all the potential human body-related variables is to take many 2T impedance measurements of two electrodes tied to a human volunteer over the course of months. Why 2T instead of 4T? There is one main reason: the contact impedance of the electrodes is the desired quantity to measure! A 4T measurement would remove the contact impedance and only measure the body impedance (significantly less than the electrode impedances).

Once a human subject is attached to electrodes and an impedance meter, the next step is to measure the impedance phase and magnitude across a large frequency spectrum (near DC to MHz). *In vivo* spectroscopy will ensure that the expected pole and zero behavior during real-life use will be recorded and an electrode model can be adequately fit. Skin phantoms can also be used, but without sweat, oil, proper skin texture, or the underlying bone structure, any extracted information would be of limited utility because it would have little relation to a real-life system. When possible, take *in vivo* (in/on body) measurements!

Electrode Half Cell Potential

The electrode DC offset (EDO) of the electrodes relative to the reference (two electrode half cell potentials added together) has implications on the required neural recording input range or offset cancellation range. Maintaining a low EDO is desirable to keep the recording frontend in its linear range and out of saturation. As with all of these quantities, there are multiple ways to measure an electrode's half-cell potential. In electrochemistry, the most common method involves placing the test electrode in an electrolyte solution and completing the circuit with a reference platinum electrode in acid (that has a known half cell voltage). While this method works, like other *in vitro* methods, it does not transfer to in-vivo systems directly.

Another, more practical method involves taking two test electrodes and attaching them to a grounded human subject as in Figure 2.17. Assuming no external bias is provided to these electrodes (signified by 0 Amps being injected onto the body), then the measured voltage across electrodes 1 and 2 would be the total electrode DC offset. It is important to note that the quantity measured across both electrodes is not an individual electrode half cell potential. Instead, it is the difference between the two half-cell potentials. It is also more helpful because it details the possible range of DC offsets for particular electrodes on a human subject. In addition, this measurement includes the influence of sweat, surface grime, oils, and day-to-day skin moisture differences. Assuming the skin around E1 and E2 is similar, the EDO should be normally distributed around zero volts. Due to the laundry list of environmental differences occurring around and on a user's skin, this EDO may vary across users or time. That is why it is essential to measure some proxy for half cell potentials

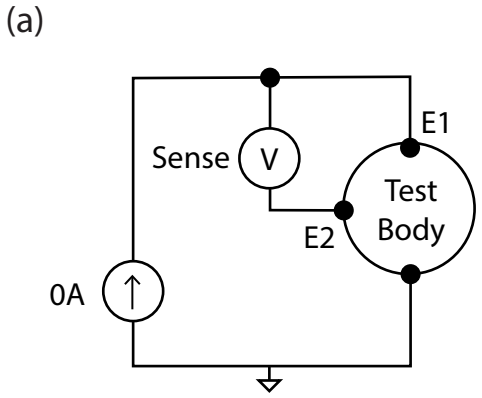

Figure 2.17: *In vitro* half-cell test setup. Two electrodes, E1 and E2, can be applied to a surface or placed in a beaker filled with conductive electrolyte (test body). The voltage measured between the two electrodes is the EDO assuming no external bias or current is being applied to the test setup.

on a body across multiple individuals across multiple days/weeks/months. There is no other way to get this data than to do it in as realistic a scenario as possible.

Lifetime Testing

When investigating how device performance can change over its lifetime, one should repeat any/all of these characterizations at set intervals. Measuring a device's resistivity every week over six months would showcase how fast/slow the material oxidizes and how that may affect performance with use. Furthermore, a flexible electrode may stiffen or become more flexible with repeated use. Understanding how its spring constant (at a constant strain) changes may provide insight into how the device fatigues over time. These characterizations should be similarly controlled at each time point (do not change experimental methods halfway through a lifetime test) to limit the amount of confounding variables and provide clear insight.

Chapter 3

User-Generic Earpiece Design and Fabrication

As discussed in Chapter 2, there is a wide array of electrode designs and fabrication techniques. Concepts of concavity, overhangs, and skin contact through hair and oil also impact earpiece design, and as a result, not all earpiece designs can be easily manufactured. As a result, when designing a user-generic earpiece that can be used across most of the human population, it is paramount to keep manufacturability in mind. Furthermore, the recording site's constraints should be leveraged to maximize electrode skin contact and mechanical stability. For example, the ear canal is a mostly closed surface. As a result, electrodes can apply opposite normal forces off each other to apply gentle pressure on the canal's surface. However, utilizing domain-specific knowledge such as ear canal morphology is not enough. It is equally essential to find high-value recording locations inside the ear. Thus the earpiece design process must first begin with an understanding of electrode placement. Only after an electrode scheme is determined can a user-generic earpiece be designed to maximize electrode contact and comfort in a user-generic fashion.

3.1 Earpiece Design

To methodically determine unique recording sites in the ear, preliminary measurements were taken with an initial user-specific earpiece designed to provide consistent electrode-skin contact. The initial earpiece was designed based on a high-resolution scan of a subject's ear (Fig. 3.1b) and made to support a relatively dense electrode array. Ten 12 mm^2 electrodes were spaced evenly across the earpiece's surface to record along much of the outer ear (Fig. 3.1c). After construction, this custom earpiece was used to record eye blinks and alpha band modulations in a single subject (see Chapter 5 for signal background).

Pearson correlation coefficients were calculated for each electrode pair from four eye blink measurement trials and four alpha band modulation experiments and used to highlight which electrodes recorded the most unique data. Average correlation coefficients for each electrode

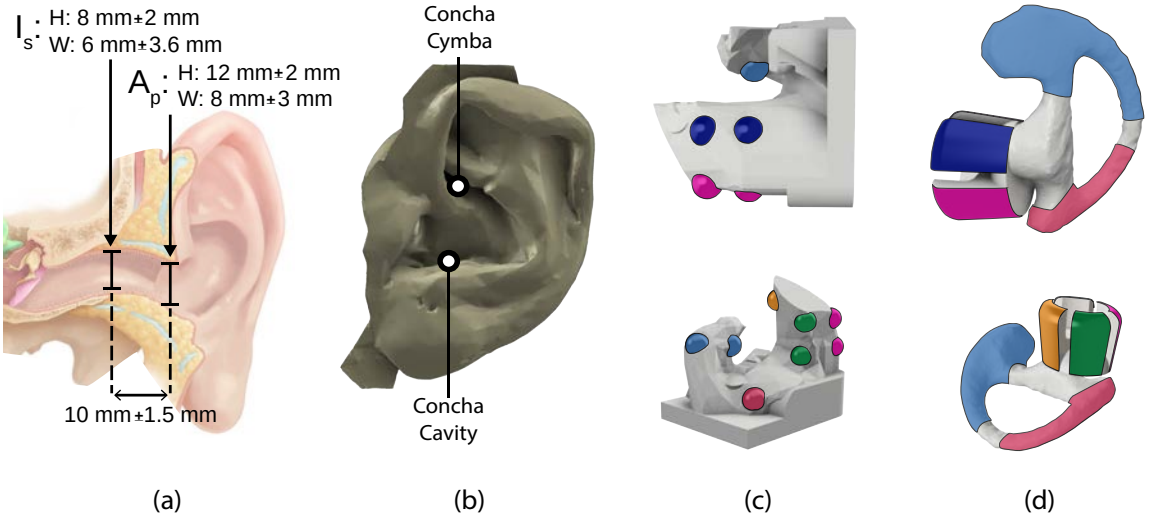

Figure 3.1: (a) Ear canal with average measurements from [69] [70]. (b) High-resolution subject ear scan. (c) Earpiece custom-fit for (b) shown from two angles. The same color highlights correlated electrode pairs. (d) User-generic earpiece. Electrodes highlighted with corresponding group colors from c.

pair were then used as criteria for combining electrodes into a larger electrode (to reduce ESI without sacrificing unique recording sites). Electrode pairs that were >70% linearly correlated were grouped (highlighted by the same colors in Fig. 3.1c and d). Electrodes with the lowest degree of correlation with all other electrodes (<20%) were selected as reference locations. The final design includes four in-ear electrodes (for sensing) and two out-ear electrodes (for referencing). Consolidating ten small electrodes into six larger electrodes reduces ESI, noise, and interference and promotes increased electrode-skin contact across different ear canal shapes. Furthermore, this electrode arrangement delivers greater flexibility with two large electrodes that can serve as a reference on the outer ear. Due to user-to-user variation in ear morphology, the concha cymba may perform as a better reference location for an individual than the concha cavity (or vice versa).

After selecting recording sites, the next step was to determine a user-generic earpiece design based on anatomical measurements of the average ear canal. The three most crucial ear dimensions (the aperture, isthmus, and length highlighted in Fig. 3.1a) are reported to have a normal distribution across large human populations [69]. Thus, three generic earpiece candidates were designed with these average measurements and then scaled up and down by a standard deviation to make small and large earpieces. Multiple individuals tested each design for comfort and physical stability during three activities: walking, running, and jumping. The most comfortable user-reported design that remained secure in the ear

through all movements was selected as the final structure (Fig. 3.1d). The selected earpiece leverages four outward cantilevers to apply independent pressure on each in-ear electrode. This compliance, coupled with the enlarged electrode area, minimizes ESI across multiple users relative to the custom-fitted earpiece and other dry electrode-based earpieces [20]. The final design has four 60 mm^2 electrodes around the ear canal's aperture and two larger, 4 cm^2 electrodes along the outer ear's concha cymba and concha cavity (Fig. 3.1a & b). After scaling the average earpiece to make small, medium, and large versions of the user-generic earpiece, it is expected that upwards of 90% of users can achieve a comfortable fit. This has so far been corroborated by small-scale testing where 29 out of 30 subjects have reported comfortable fit and acceptable electrode skin impedance for at least one earpiece size.

3.2 Earpiece Manufacturing

After selecting electrode locations and a user-generic earpiece design, the final step is constructing the earpiece. To explore different manufacturing techniques and their respective scalability and iterability, two main fabrication methods were investigated: Vacuum forming and 3D printing.

Vacuum Forming and Spray Coating

Vacuum forming is commonly used in the mass production of plastics for everything from toys to automotive parts. Forming techniques allow the reuse of a small number of molds to create many parts and thus have minimal tooling costs for large-scale manufacture. The drawback is that certain features such as overhangs and cantilevers are not easily formed, and the material thickness will vary depending on the curvature of the design.

Constructing an earpiece with vacuum forming and spray coating is generally straightforward. Following the process steps listed in chapter 2 yields results shown in Figure 3.2. A mold (Figure 3.2a) is 3D printed and vacuum formed with polycarbonate (Figure 3.2b). Once the hollow form is liberated (Figure 3.2c) and spray-coated (Figure 3.2d), the resultant earpiece has cantilevered electrodes that can apply gentle pressure on the ear canal (Figure 3.2e). It is important to note that small overhangs are sometimes possible to construct if the mold is appropriately oriented but will require cutting the shell off the mold.

As alluded to earlier, this is not entirely pain-free. Due to the cantilevered nature of the in-ear electrodes and some inherent overhangs associated with a 3D, partially enclosed surface, the earpiece must be physically cut out of the excess polycarbonate. Through experience, the easiest way to do this was by using a Dremel to remove material from the inside of the ear canal portion and cutting the entire earpiece in half from the cymba (y) to the concha bowl electrode (c), bisecting the cantilevered electrodes. This step required skilled labor to not further stretch the already thinned polycarbonate shell (i.e., cutting and pulling off the mold risks stretching the surface).

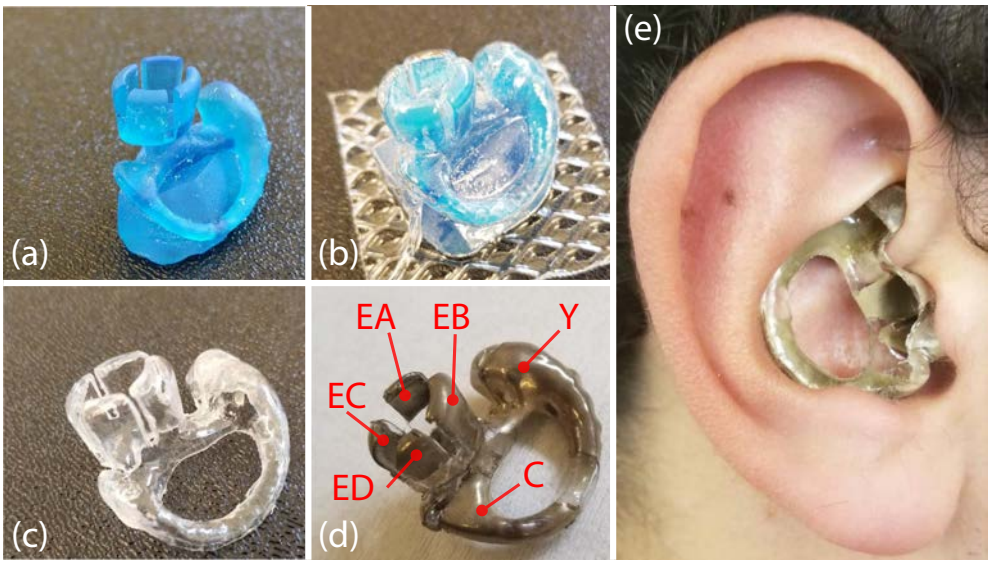

Figure 3.2: Earpiece construction and in-ear fit. (a) 3D printed master mold. (b) Thermoformed earpiece. (c) Liberated earpiece. (d) Spray-coated earpiece. (e) Earpiece in-ear fit

Furthermore, the silver spray-coated electrodes have excellent conductivity and cohesive surfaces (when sprayed after forming) but cannot be soldered. As a result, all interconnects must be made with conductive epoxies, which are not as mechanically stable and robust as traditional metal solder. Lastly, due to the characteristics of a thin thermoformed plastic substrate, the earpiece will have some compliance, but will not be compliant across multiple axes at the same time. This rigidity makes vacuum-formed earpieces less comfortable from a user perspective.

Regardless, vacuum forming allows for the construction of multiple earpieces in parallel if the vacuum former has a large enough bed. This means multiple different designs can be tested out at the same time and constructed in tandem. Thus for small-scale testing, vacuum forming can enable rapid iteration on different designs.

3D Printing and Electroless Plating

3D printing is less practical for large-scale manufacturing but offers greater freedom for design exploration. Complex structures can be easily printed without regard to their concavity. Furthermore, multiple different materials can be used. The tricky part is the metallization of the surface, which is where electroless plating comes in.

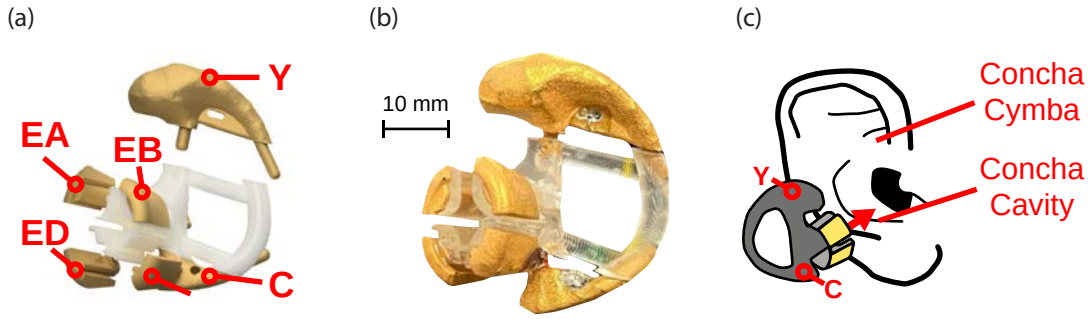

Figure 3.3: Photograph examples of 3D printed and plated earpieces

First, it is important to highlight which parts of the earpiece need to be rigid and metalized (golden electrodes in Figure 3.3a). These pieces should be designed with interconnects and material stability in mind. Make sure there are loops for wires to attach to and rails to help orient and fit the electrodes on the skeleton. The (to be) metalized parts should be printed using a rigid material, while the skeleton should be printed with an elastic material. This ensures that the earpiece will be compliant across any axis, unlike the polycarbonate shell that could only bend in the direction of the lever in question. This improves comfort and ideally improves earpiece fit and stability inside the ear. The rigid pieces should then be processed and plated according to the Au plating guide in chapter 2.6.

After plating, the earpiece can be assembled like legos (Figure 3.3b). Wires can be soldered directly to the electrodes (like conventional electronics), making assembly significantly faster than the vacuum forming case. It is recommended to cover any solder joint with clear epoxy to passivate the surface and provide further mechanical stability. Furthermore, using thin gauge wires (30 AWG) will reduce wire stiffness and improve the earpiece's useability. While less scalable from a commercial perspective, 3D printing manifold earpieces with soft and rigid materials significantly improve user comfort, ease of assembly, and design space freedom.

3.3 Earpiece Characterization

Electrode Skin Impedance

The single most crucial comparison metric between two functional electrophysiological are their electrode-skin interface impedances, and these earpieces perform comparably. Figure 3.4 plots the average ESI magnitude and phase of the vacuum-formed devices, spray-coated earpieces, and the 3D printed gold plated earpieces and a clinical gold cup wet electrodes. No skin preparation or cleaning was performed before measurements to simulate real-world conditions. It is important to note that the wet electrodes (Au cup - hydrogel) are smaller than the dry electrodes, and this smaller electrode surface area results in a higher measured ESI than is expected for a wet vs dry electrode comparison

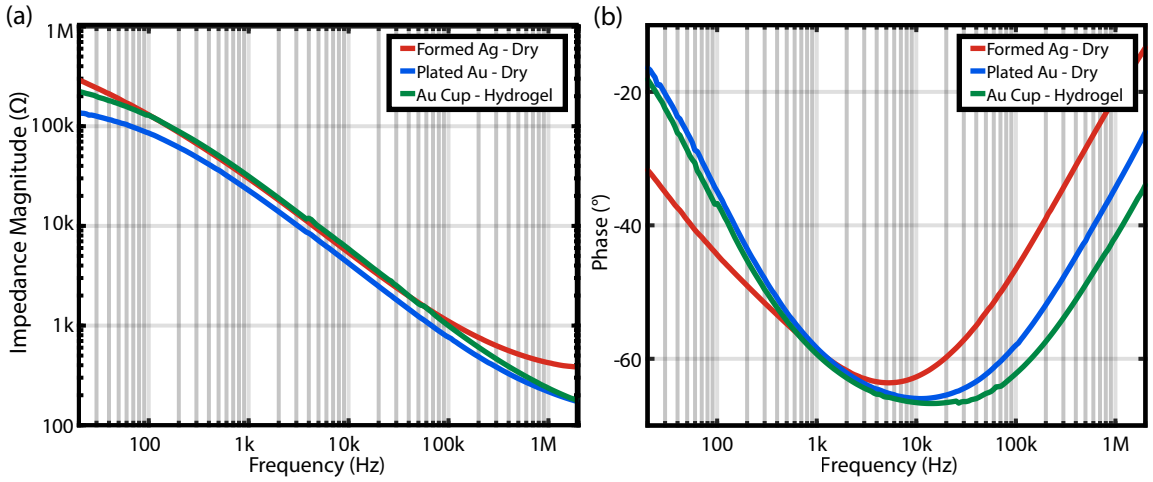

Figure 3.4: Average (a) magnitude and (b) phase comparison of various vacuum-formed spray coated dry electrodes (n=100), 3D-printed gold plated dry electrodes (n=25), and gold cup electrodes with hydrogel (n=10).

An important takeaway is that the 3D printed and gold plated earpieces exhibit a lower ESI magnitude than their vacuum-formed counterparts. This is likely due to their increased surface roughness (greater effective surface area) and compliance. More compliant earpieces imply that the electrodes will be making better contact in the ear canal - thus further increasing the geometric contact area.

Electrode DC Offset

Due to the sheer number of measurements taken with the initial user-generic earpieces, there is significantly more EDO data for the vacuum-formed and silver spray-coated earpieces

than for the 3D printed, gold-plated earpieces. Figure 3.5 plots a histogram of all EDO measurements taken with the silver spray-coated earpieces.

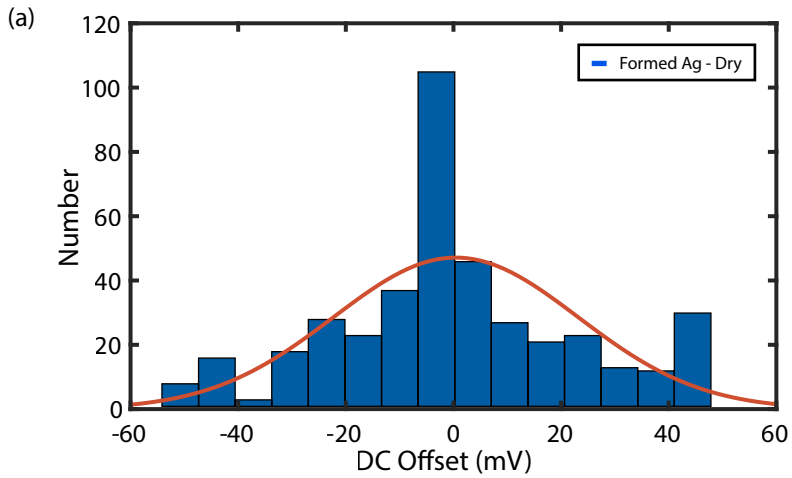

Figure 3.5: Histogram plotting electrode DC offset (EDO) measurements across three users wearing a thermoformed Ag coated earpiece (n=395)

The formed Ag EDO measurements are normally distributed (n=395), with a mean and standard deviation of 0.444 mV and 22.7 mV, respectively. The minimum and maximum EDO values are respectively -49.6 mV and 45.1 mV.

There are significantly fewer measurements (n=10) for the Plated Au earpieces, and thus there is considerably more noise, and no clear distribution exists yet. The average value is one mV with a min and max EDO of -30 mV and 60 mV, respectively.

Compliance and User Comfort

Each electrode was strained with a force gauge to measure the in-ear electrodes' (EA, EB, EC, and ED) compliance on the vacuum-formed earpiece. The average electrode spring constant was 171 N/m with a standard deviation of 5 N/m (n=50). This inherent spring allows each electrode to apply approximately 90 kPa (at 50% strain). This pressure is on the order of polyurethane foam (at 50% strain).

The 3D printed and gold-plated earpiece had a softer spring constant of 130 N/m with a standard deviation of 10 N/m (n=50), making it softer than the vacuum-formed earpiece. Furthermore, the entire gold earpiece's skeleton had the same spring constant, unlike the significantly stiffer polycarbonate shell.

This difference in spring constant and body compliance was further supported by user-reported comfort. Of the ten users that tried both types of earpieces, all ten reported that the 3D printed and gold plated earpieces were significantly more comfortable.

3.4 Improvement Strategies

While the 3D-printed gold-plated earpieces offer more comfort and lower impedance, this design could be further optimized based on experimental and physiological measurements. As was initially done, the earpiece can go through another round of optimization based on experimental and physiological measurements. It is also entirely possible that different applications may require other electrode counts, areas, and shapes. Subsequent chapters will describe the remaining parts of a wearable system, namely the readout electronics and the application-focused software. Through these chapters, emphasis will be placed on the system-level analysis that can be used to iterate on not just system specifications but individual earpiece and electronic design parameters. These iterations will ideally lead to simpler, more robust, and more comfortable wearables that can achieve greater adoption.

Chapter 4

Recording Readout and System

A basic recording system is generally composed of three blocks. An amplifier will buffer and amplify signals from the sensor, an analog-to-digital converter will transform the analog EEG signal to a digital signal that can be efficiently handled by modern computing hardware, and lastly, a digital backend that can post-process/store/transfer data from the readout to the outside world. These blocks are drawn out in Figure 4.1 with an electrode model, an amplifier with input resistance R_{in}, an ADC with output bitwidth N, and an antenna. The recording system does not include the electrode and is highlighted with a blue box.

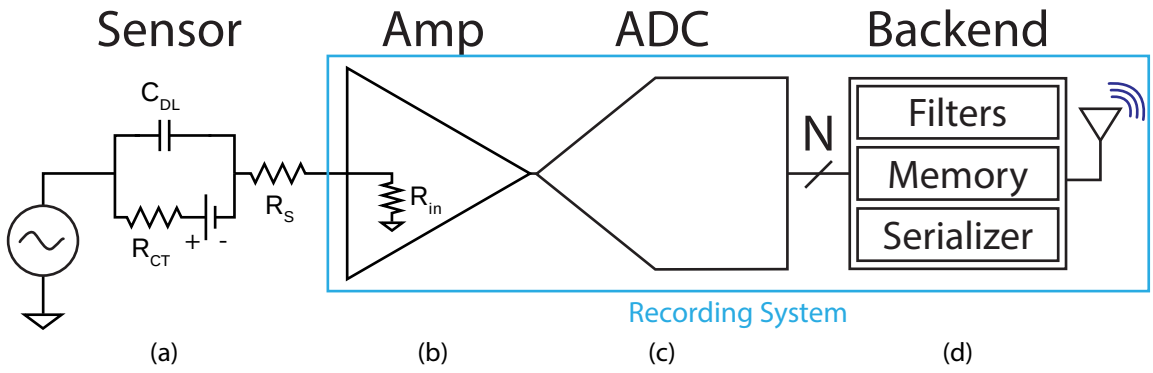

Figure 4.1: Generic, single channel, readout schematic from electrodes to a wireless Interface with an N bit ADC output. Some chip architectures will combine low-noise amplifiers (LNAs) and ADCs or omit backends, but the same system concepts are applicable.

While incredibly important to the overall system functionality, this report is not focused on the analog design techniques required to design and build high-performance circuits. Rather, the goal of this chapter is to discuss general parameters and information important to the system design of Ear EEG hearables. After each block and the most important

specifications are introduced, subsections will touch on how to understand certain system characteristics such as system input impedance, noise, and resolution. Lastly, there will be a brief summary and synthesis of the experiments required to iterate and tailor a system to a specific application.

4.1 Basic Blocks

The readout amplifier, ADC, and backend all play important roles in recording EEG. Different research systems may remove or combine these blocks, but their functionality is always required for a complete neural recording system to function. Below are high-level descriptions of each block with their pertinent performance metrics.

The initial amplifier, amp, is usually the first block in a readout chain and is generally defined by its input impedance, input referred noise, input range, gain, and bandwidth. The amplifier's job is to buffer as much of the EEG signal as possible, amplify it by some gain, and pass it to any downstream digitization blocks. An amplifier's input range and bandwidth act as limits on what input signals can be buffered. For example, a 10 mV (in amplitude) signal would likely saturate if the amplifier only has an input range of 4 mV. While one might argue that EEG signals are usually $< 1mV$ in amplitude and thus amplifiers with small input ranges are sufficient, keep in mind that motion artifacts or electrode DC offsets could be well above 5-10 mV. Furthermore, if an ear-worn wearable is meant to record EEG and muscle movements, the amplifier must have a bandwidth that encompasses both the EEG and EMG bands; otherwise, those signals will be filtered out.

After the amplification stage, an ADC is required to digitize the recorded signal. At this point, input impedance should no longer be a concern because the ADC should not directly load the electrodes and the most pertinent performance metrics are the signal-to-noise-and-distortion-ratio (SNDR), input full-scale, and output sample rate. SNDR and bitwidth ultimately determine the linearity and resolution of the digitized signal and are some of the most important system specifications. Furthermore, SNDR also determines the effective number of bits an ADC can output. There are many commercial ADCs that may be marketed as 22bit ADCs but they only have 16-18 bits worth of SNDR for a given conversion time. The rest of the bits may be noise or at least heavily distorted. In addition to linearity and resolution, a system's SNDR will also be impacted by the total readout chain's noise: common mode and supply noise in the amplifier and ADC, ADC quantization noise, thermal noise, shot noise, etc. The ADC's input full-scale is analogous to the amplifier's input range. The output sample rate pertains to more than just output bandwidth because it also determines the data throughput and sets requirements for the wired/wireless link.

The digital backend is sometimes integrated onto the same substrate as the amplifier and ADC (as is the case in section 4.6), other times it is a separate block on another board. Regardless, data from an ADC must be processed and transferred to other computing resources to actually be used. The digital backend is generally the interface facilitating this transfer but may vary wildly depending on the specific system implementation. Some

systems call for a fully integrated system on chip with digitization, signal processing, and classification together while others will be able to transfer all the data off-chip and process it on a microcontroller or cloud devices. Usually, power and application requirements will dictate which implementation to pursue. For the purposes of this report, we will assume the backend is solely responsible for transferring any data to a wireless link for offline processing.

It is important to note that some chip architectures will combine low-noise amplifiers (LNAs) and ADCs or omit digital backends. Many different architectures exist [71] [72] [66] but the system concepts mentioned above are always applicable. When determining the requirements of an entire system, input impedance, SNDR, sample rate, etc. will all be important performance metrics for comparing different architectures and devices.

4.2 Electrodes and Amplifier Input Impedance

In its simplest terms, the amplifier input node is a voltage divider with the human body as the source. The voltage at the inputs of the differential amplifier in Figure 4.2 can be written as:

$$V_{in+} = \frac{1}{2} V_{Head} \frac{R_{in+}}{R_{in+} + Z_{sens}}, \tag{4.1}$$

$$V_{in-} = -\frac{1}{2} V_{Head} \frac{R_{in-}}{R_{in-} + Z_{ref}}, \tag{4.2}$$

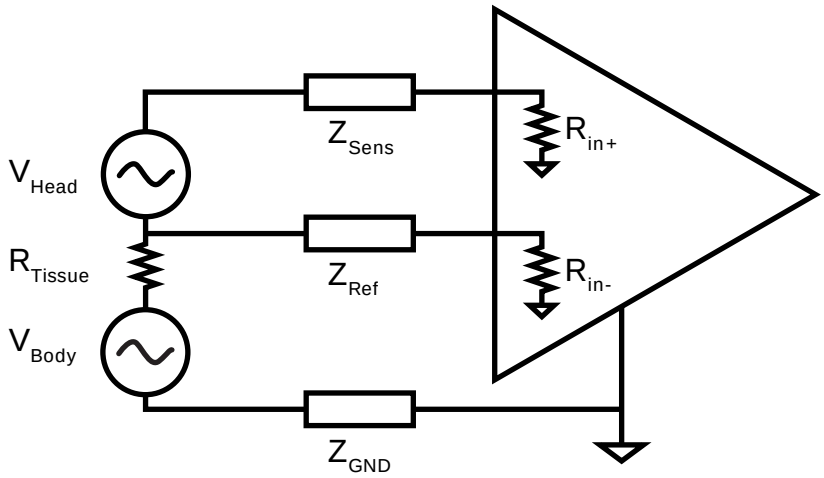

Figure 4.2: Close-up of Two Electrodes Plugged into a Differential Amplifier

Then the total input voltage would be:

$$V_{in} = V_{in+} - V_{in-}. \tag{4.3}$$

These voltage dividers showcase just how important it is to have a high input impedance. If the electrode impedances (Z_{sens} and Z_{ref}) are the same as the input impedance then 50% of the signal amplitude is immediately lost. In addition, the differential nature of these measurements makes the system immune to common mode signals if and only if there is no mismatch between the two input paths. If the electrode impedances are different, common mode signals such as (50/60 Hz) interference from the power lines will immediately be converted to differential signals and hurt the signal's SNR. It is largely impossible to control the mismatch of the two different electrodes since humans do not sweat the same nor are we physically shaped the exact same. The best designers can do is ensure that all electrodes have the same geometric surface area and maximize the amplifier's input impedance. If the amplifier's input impedance is significantly larger than the electrode impedances, any stationary electrode-related common-mode to differential-mode conversion is minimized.

Excessive interference can also be mitigated by active driven right leg (DRL) and passive grounding techniques. DRL techniques focus on measuring background signals on the human body (e.g. 60 Hz from power lines), inverting them, and then driving the body with a destructively interfering signal [73]. Historically, ECG systems would place this DRL electrode on the right leg which is where the name comes from. This technique is also sometimes referred to as an 'active grounding' method. Similar effects can be achieved by attaching a low-impedance electrode tied to the system ground. This will sink some interfering current. Ultimately, DRL and passive grounding techniques hope to reduce the total power of interferers, but they cannot completely eliminate them.

4.3 Input Referred Noise Sources

To robustly record in-ear EEG, the total system noise floor must be low enough to provide sufficiently high SNR for all features of interest across a wide population of users. State of the art typically quantifies the noise arising from the electrode and the recording electronics independently. However, to analyze the total input noise, all components and their interactions must be considered, as shown in Fig. 4.3.

The electrode noise can be modeled as the thermal noise contributions of R_s (v^2_{n,R_s}) and R_{ct} ($i^2_{n,R_{ct}}$) [44]. Typically, noise from recording electronics is modeled as a combination of uncorrelated input-referred voltage and current noise. However, these two noise sources can be correlated, since they may be generated by the same devices [74]. In this case, the noise can be modeled as a combination of the input-referred voltage noise $v_{n,a}$ and current noise $i_{n,a}$ as well as a correlated noise contribution [75]. S_{vv} and S_{ii} are the power spectral densities (PSD) of the voltage and current noise respectively while S_{vi} describes the PSD of the correlated component of the noise. The total noise S_{nn} referred to as the input of the

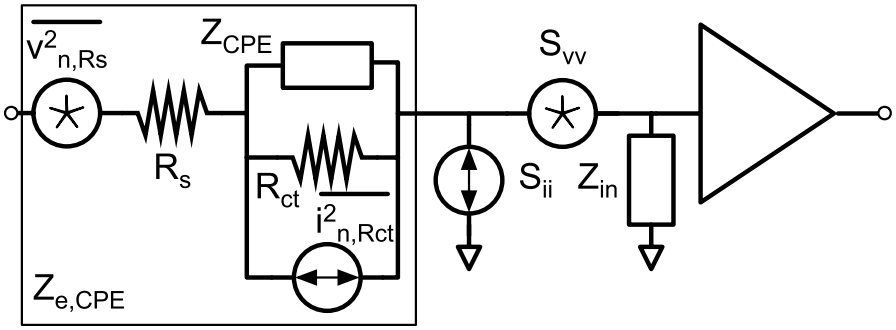

Figure 4.3: Noise sources at input of readout frontend: electrode noise $v_{n,Rs}$ and $i_{n,Rct}$, input-referred frontend voltage noise $v_{n,a}$, and input-referred frontend current noise $i_{n,a}$.

electrode can be described by

$$\begin{aligned} S_{\mathrm{nn}} = & \overline{v^2_{\mathrm{n},R_\mathrm{s}}} \\ & + \overline{i^2_{\mathrm{n},R_\mathrm{ct}}} \left| \frac{R_\mathrm{ct}||Z_\mathrm{CPE}}{(R_\mathrm{ct}||Z_\mathrm{CPE}) + R_\mathrm{s} + Z_\mathrm{in}} \cdot (Z_\mathrm{in} + Z_\mathrm{e,CPE}) \right|^2 \\ & + S_\mathrm{vv} + S_\mathrm{ii} \left| Z_\mathrm{e,CPE} \right|^2 \\ & + 2\mathrm{Re} \left\{ S_\mathrm{vi} \frac{Z_\mathrm{in} Z_\mathrm{e,CPE}}{Z_\mathrm{in} + Z_\mathrm{e,CPE}} \right\}. \end{aligned} \tag{4.4}$$

Due to the small value of R_s (<1 kΩ) relative to R_{ct} and $Z_{e,CPE}$ (both >100 kΩ), the contribution of the first term is negligible. The contribution of S_{ii} and S_{vi} are often ignored since the typical wet electrode impedance is sufficiently small such that other noise sources dominate [44][76]. However, in the context of dry electrode recording, it cannot be ignored due to the high electrode impedance; therefore the current noise contribution can increase the noise floor and degrade SNR. A more detailed discussion performed by Justin Doong can be found in [35]. Other electrode noise analyses can be found in [77] and [78].

There is another noise source to consider, power line interference. While band-limited, interference from the power lines (50/60 Hz) can be large enough to saturate the front end. Common mitigation techniques include driven right leg or 'active' grounds that drive the body with a destructively interfering 50/60 Hz signal [48]. A simpler technique involves having a low impedance ground electrode on the body (passive ground) [79]. Grounding the body near the measurement sites greatly reduces the 50/60 Hz signal power, straightforward as that. No fancy techniques are required.

4.4 Input Range and Coupling

As mentioned earlier, the readout's input range must be large enough to resolve signals of interest even in the presence of interferers. When recording signals from inside the ear, motion and muscle-related artifacts can reach amplitudes $> 5mV$ while EEG signals are $< 10\mu V$. Figure 4.4 shows typical chewing-related artifacts. A subject was tasked to bite out of an apple every 20 seconds. Each initial bite is at least $3mV$ and the chewing-related signals are limited to $< 10Hz$ as shown by the time-frequency spectrogram in Figure 4.4. In addition to motion-related artifacts, electrode DC offsets can also saturate a front end.

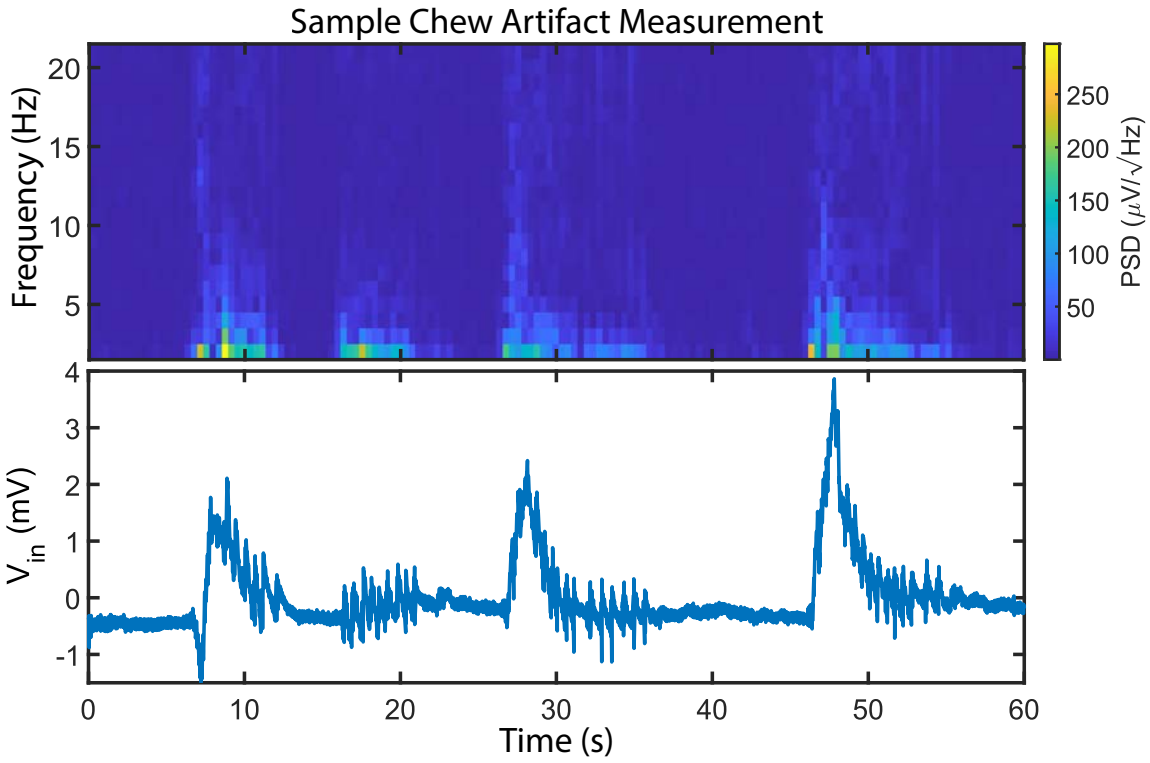

Figure 4.4: Single channel in-ear measurement of a subject chewing an apple

Figure 3.5 showcases the possible EDO values (-50 to 50 mV) that can appear at the input.

There are multiple techniques for mitigating these large low-frequency interferers. Many systems are AC coupled, meaning there is a capacitor between the electrodes and the front-end amplifier. As a side note, DC coupled systems have no capacitors between the electrodes and front-end. AC coupling the front end blocks DC signals, such as EDO, from saturating the signal. Unfortunately, AC coupling capacitors introduce large settling times (so any step-like signal will have a long tail) at the input. Furthermore, AC coupling will not eliminate signals above DC such as chewing or walking-related artifacts. In those cases, it is best to

characterize the largest possible signal magnitude and design a system that can handle it with feedback circuits or amplifiers with the inherently large, rail-to-rail operation. This way EDO can also be digitized and used to periodically characterize electrode behavior.

4.5 Wired vs. Wireless Interfaces

The last bit of the readout system will be its connection to the outside world. Oftentimes, initial setup and calibration are easier and more straightforward when done over a wired interface. In the most basic form, this wired interface can be a digital, generic serial peripheral interface (SPI). The benefit of using a SPI interface is that most microcontrollers and development kits support SPI, and this will provide a lot of freedom during the initial system bring-up. For example, common microcontrollers such as the Nordic NRF52 series and Atmel ATmega series all come with at least one if not on chip or onboard multiple SPI controllers.

After the initial bring-up, wired interfaces (power, timing, SPI, etc.) may limit what type of on-body measurements can be performed. This is purely from a practical and experimental perspective. If there is too much support hardware, it will be hard to have a user comfortably perform a task while hooked up to the readout system. This is where battery-powered, compactly integrated systems with wireless interfaces come in. Not only are fully wireless systems much simpler to use in experiments, but they are also less mechanically temperamental. Motion on a serial interface or power cable can cause artifacts to appear in neural recordings due to long lever arms and the torque wires can place on sensors. The downside to a wireless interface is that battery life becomes a concern and extra hardware will be required for radio or Bluetooth links. Luckily, this functionality exists in commercial ICs and they can be easily added to a PCB.

The primary hurdles for implementing wireless interfaces are the setup time, data throughput, and power consumption. Wireless interfaces tend to have higher development time due to their increased complexity. However, commercial vendors have significantly streamlined the setup process with extensively debugged application programming interfaces (APIs) and easy-to-use development kits. Nordic Semiconductor is one such example, their NRF5 Bluetooth platform has extensive documentation, examples, and technical support to aid in the bring up of a Bluetooth system. Other companies such as Espressif and GreekCreit have implemented similar solutions for WiFi modules and RF transmitters and receivers.

Wireless interface max data throughput and the corresponding power consumption are intertwined. Typically, higher throughputs require faster transfer rates which consume more power. Depending on the channel count and data bitwidth, BLE can provide an elegant solution. The Bluetooth 5 standard can theoretically achieve data rates up to 2 Mbps. In reality, when accounting for distance and added packet security, it will likely be 1.75 Mbps. Wifi and RF links can achieve much higher throughputs at the cost of power, which is the real prohibitive barrier for wearable devices.

4.6 Don't Reinvent the Wheel: WANDmini

When it is possible, it is always best to use tried and true hardware that has been calibrated, tested, meets desired system performance and is low-power. One such system is a miniature, wireless, artifact-free neuromodulation device (WANDmini), a low-profile, custom neural recording system that streams recorded data over Bluetooth Low Energy (BLE) to a base station connected to a laptop. WANDmini is derived from a previous design for a wireless, artifact-free neuromodulation device (WAND) [80], reduced to a form factor of $2.5 \times 2.5\ cm^2$, and embedded with custom firmware (Figure 4.5). This system was augmented and used to perform Ear EEG recordings with the help of Andy Zhou and Fred Burghardt [35].

Recording and digitization are performed by a custom neuromodulation IC [72] (NMIC, Cortera Neurotechnologies, Inc.) integrated with 64 digitizing frontends, thereby expandable to recording applications with higher electrode counts. NMIC and WANDmini specifications are listed in Table 4.1, and the WANDmini system block diagram is shown in Figure 4.6.

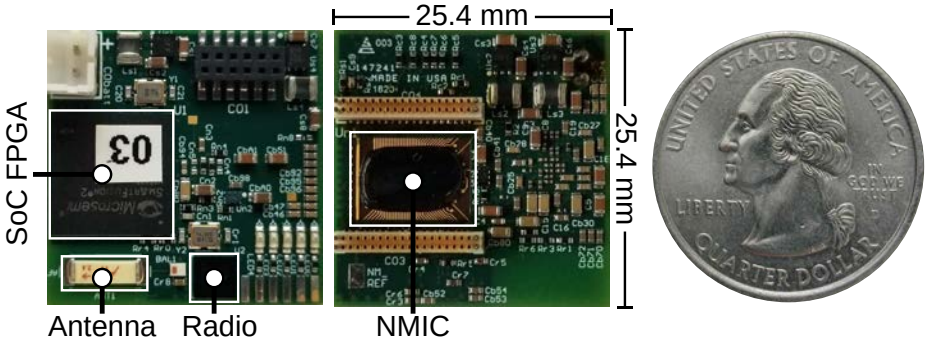

Figure 4.5: WANDmini neural recording module beside scale quarter.

The NMIC is an ideal starter system due to its low power and high dynamic range, supporting a 100–400 mV input range with a flat input-referred noise voltage spectrum of $70\ \text{nV}/\sqrt{\text{Hz}}$. The analog-to-digital converters (ADCs) have a resolution of 15 bits and sample at 1 kS/s, providing sufficient resolution and bandwidth for EEG signals. The wide linear input range can accommodate the large electrode dc offsets (Fig. 3.5) and provides robustness to interference. The total harmonic distortion (THD) with a full-scale input is 0.012%, maintaining linearity in the presence of large interferers [80]. The input range is expandable to 400 mVpp to accommodate larger electrode offsets at the expense of quantization noise. The front-ends achieve this large range with a mixed-signal architecture that includes the ADC into the feedback loop to reduce the required gain and signal swings. Many systems handle EDO by simply AC coupling, but recording EDO provides additional information on electrode performance and eliminates the long settling times associated with AC coupling capacitors. The NMIC also has stimulation and impedance measurement capabilities, which are not used in the Ear EEG system.

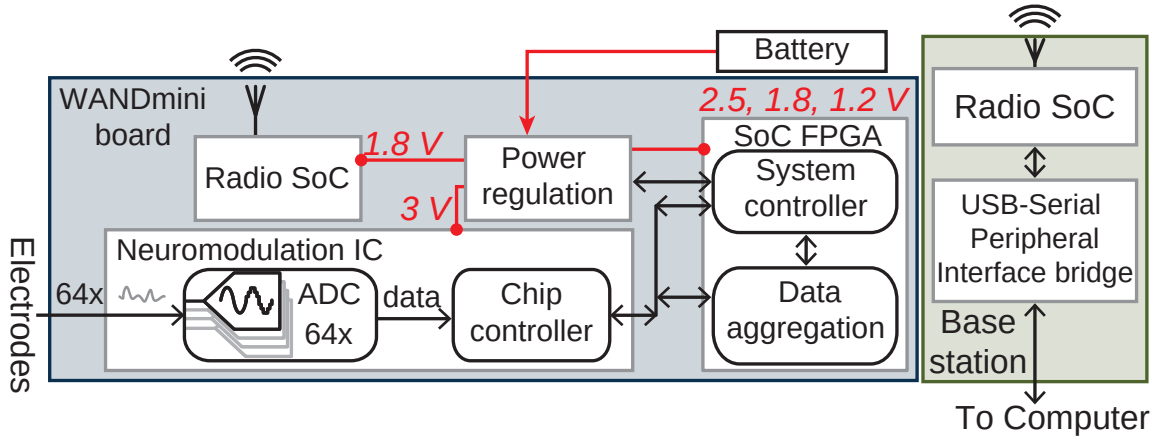

Figure 4.6: WANDmini system architecture as used in the user-generic ear EEG

Electrode models can be used to emulate electrode noise behavior in a controlled environment. As described earlier, the input current noise contributes to the overall system noise when coupled to a high impedance electrode. An input current noise of 286 fA/$\sqrt{\text{Hz}}$ and a correlated noise of 11.2 $\sqrt{\text{W/Hz}}$ were measured for the NMIC recording channel. With the electrodes present, the low-frequency noise spectrum increased to about 270 nV/$\sqrt{\text{Hz}}$ and resulted in an integrated noise of 4.65 $\mu Vrms$, sufficient to record ear EEG with SNR >15 dB. Since the electrodes contribute about 90 nV/$\sqrt{\text{Hz}}$, a custom design with lower current noise terms could decrease the total integrated noise by up to 2 $\mu Vrms$. The electrodes were modeled as $Z_{e,CAP}$ with parameters $R_s = 1k\Omega$, $R_{ct} = 270k\Omega$, and $C_{dl} = 2.4nF$.

Table 4.1: NMIC and WANDmini Specifications

NMIC	
Recording Channels Max/Used	64/5
Input Range	100 mVpp
Input-Ref Voltage Noise	70 nV/$\sqrt{\text{Hz}}$
Input-Ref Current Noise	286 fA/$\sqrt{\text{Hz}}$
THD	0.012% (100 mVpp input)
ADC Resolution	15 bits
ADC Sample Rate	1 kS/s
Power	700 μW
WANDmini	
Wireless Data Rate	2 Mbps
Board Dimensions	25.4 mm $\times$ 25.4 mm
Weight (w/o battery)	3.8 g
Supply Voltage	3.5 V
Power	46 mW

4.7 Potential Improvement

Research platforms are generally optimized for flexibility and high performance in order to not limit users to narrow use-cases. This often comes at the cost of complexity and power. Once these platforms have been used to highlight specific applications, future systems can be optimized to reduce channel count and target the specifically pertinent bandwidths and signal amplitudes. These simplifications can lead to reduced power consumption in the readout as well as simpler interfaces that may not need to stream data all the time. For example, classification or post-processing may be able to be done on-chip and the result streamed out as opposed to sending all the raw channel data. These power-reduction improvements increase battery life and better enable the periodic use of these wearable devices.

Another improvement strategy is to implement new functionality that may enable new analysis or exploration. For example, there has been growing interest in incorporating simultaneous impedance measurements onto EEG recording front ends [71]. Due to the capacitive nature of dry electrodes across 1 Hz - 2 kHz (Figure 2.7), an EEG+ESI recording front end could monitor ESI variations outside the EEG band and correlate them to ESI changes in the EEG band. Not only does this allow the grading of electrode contact and EEG recording quality, but it also enables the possible removal of ESI-coupled motion artifacts from EEG measurements. In fact, this is a current area of research being led and supported by Aviral Pandey, Justin Doong, Sina Faraji Alamouti, Cem Yalcin, Meraj Ghanbari, and Professor Rikky Muller [71]. Regardless, while implementing new features such as simultaneous ESI+EEG recording may increase system complexity, the potential benefits of long-term electrode monitoring and improved recording reliability outweigh the added complexity and cost.

Chapter 5

Experiments for Recording Neural Signals

5.1 Designing Experiments

After assembling an ear EEG system, the next step is verifying its functionality with electrophysiological recordings. These signals do not need to be cortical in origin. Initially, it is simplest to record large amplitude signals, such as electrooculography (EOG) and electromyography (EMG). EOG signals are related to eye blinks and eye movement. When a person's eye moves, the electric field around the scalp is perturbed, since the lens of the human eye has a slight positive charge and the retina has a negative charge. EOG signals are often regarded as artifacts to be removed, but they can also be a very valuable benchmark and sanity check signal. EMG signals are related to muscle activation and are also seen as artifacts when recorded on the scalp. However, they are easily controllable signals and can be instrumental benchmarking signals. To activate large amplitude EMG signals, the subject needs to clench their jaw, smile, or slowly move their head.

After EOG and EMG signals are reliably recorded, neural and cortical signals are the next step. Most neural responses, aside from alpha waves (EEG between 8-13 Hz), require a specific electrode scheme or stimulus signal. The following sections will step through experimental basics, common EEG phenomena, and showcase sample measurements taken by myself, Justin Doong, and Carolyn Schwendeman [35] [81].

Recording Basics

Three guiding principles can help with repeatably recording valuable data.

1. Signal localization and stimulus should dictate how experiments are run.
2. Minimize interference and potential noise during setup.

3. Design consistent and repeatable experiment protocols with quantitative metrics for success and result comparison to the state of the art.

4. All data should be cataloged excessively and detailed session notes should be taken.

5. Take 'gold standard' recordings with known systems.

Before running an experiment, it is essential to know what signal pattern is expected. Furthermore, knowing the target signal origin will dictate a sense/reference electrode placement scheme. Ultimately, sense electrodes should be closed to the signal source while the reference should be in a relatively 'quiet' and unrelated area. This will help maximize the recorded signal-to-noise ratio. On top of that, slight differences in stimulus may evoke different response intensities. For example, white noise and pink noise auditory stimuli will evoke slightly different types of responses from the auditory cortex.

When donning wearables, it is vital to double-check everything and ensure a tidy setup. Confirm reference and ground electrode connectivity, stow any dangling wires to prevent them from getting caught or causing the subject discomfort, and make sure all experimental hardware is fully charged and powered. After the initial setup is verified, time should be dedicated to letting electrode characteristics such as impedance and EDO settle and reach stable values. After all, any sweat build-up is likely to help reduce ESI and normalize EDO. This should not take any extra time. While the subject lets the electrodes settle and reach a steady state, they can receive instructions or fill out paperwork. Then, run sanity check measurements before any large experiment. A practical example is to start a recording and ask the user to blink their eyes. This will help verify that you see a response on all critical channels, that no connectivity issues persist, and that there is no excessive 50/60 Hz line interference. Lastly, it is essential to consider the subject's comfort and ensure they have adequate privacy while undergoing these experiments.

During and after the experiment, time should be set aside to immediately take detailed notes and catalog all recorded data. This will make an enormous difference during data analysis, especially when data may be shared with outside consultants or collaborators. Not to mention that human memory is not perfect, and detailed trial notes could potentially help explain oddities or random phenomena as they arise.

Scalp vs. Hearable Recordings

This report focuses primarily on in-ear recordings, but some understanding of scalp measurement setups may prove helpful if there is a need to compare scalp and ear EEG recordings. The sole difference between the scalp and ear recordings is their electrode arrangement. Scalp measurements are often taken with the 10/20 system (Figure 5.1). Electrodes are placed over different lobes and labeled accordingly: pre-frontal (Fp), frontal (F), temporal (T), parietal (P), and occipital (O). A central electrode line (C) and Z electrodes also represent electrodes placed on the midline sagittal plane of the skull (FPz, Oz, Cz, etc.). Cz is commonly

used as a reference during recording, but often scalp measurements are re-referenced during post-processing.

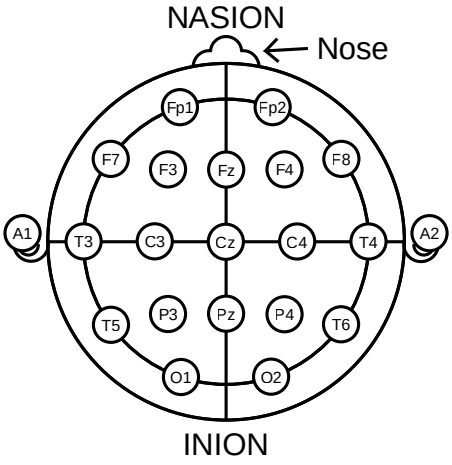

Figure 5.1: The scalp EEG 10/20 electrode scheme. Electrodes are labeled according to placement: pre-frontal (Fp), frontal (F), temporal (T), parietal (P), and occipital (O). A central electrode line (C) and Z electrodes also represent electrodes placed on the midline sagittal plane of the skull (FPz, Oz, Cz, etc.).

Re-referencing is a simple operation where two channels recorded with a shared reference are subtracted to form a new virtual channel (equation 5.3).

$$Ch_1 = Elec_1 - Elec_{ref}, \tag{5.1}$$

$$Ch_2 = Elec_2 - Elec_{ref}, \tag{5.2}$$

$$Ch_1 - CH_2 = Elec_1 - Elec_2. \tag{5.3}$$

This can help eliminate interference across sensing channels that did not appear in the reference channel.

When performing in-ear measurements, electrode locations are limited and can only be placed in the ear canals or auricle. This means existing scalp research cannot be directly re-used because EEG is so highly electrode-scheme dependant. A single reference should be placed where there is little to no EEG signal (near the ear lobe or in the concha cymba - Figure 5.2). The ground electrode should also be somewhere without an EEG signal but far from the reference. As a result, the opposite ear's concha cymba or the mastoid (behind the ear) should be grounded.

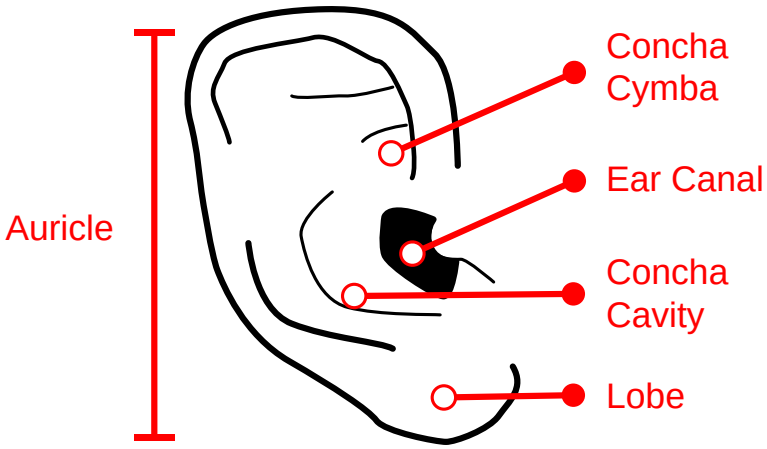

Figure 5.2: Ear Diagram for Electrode Placement. The concha cavity and cymba are ideal locations for references and grounds due to their added distance from the brain.

Controls and Brainless Experiments

When testing a new procedure, it is important to ensure that your experiment does not directly induces a response that could be mistaken for neural activity. For example, suppose your stimulus may have a mechanical component such as a tactile motor or speaker placed near the subject. In that case, extra care should be taken to ensure no mechanical coupling exists between the stimulus and the electrodes. Often times mechanical signals can vibrate an electrode, and modulate its impedance and this can appear as an EEG signal.

To confirm that there is no outside interference, brainless experiments are often performed. A setup is replicated with a mechanical phantom in place of a human subject. If an experiment showcases a 'neural' response, then, unfortunately, there is an experimental error! This is why it is vital to have clean setups when running EEG experiments!

5.2 Common EEG phenomena and confounding signals

There are several standard/known EEG experiments that people use to benchmark new systems. Some simple experiments can elicit spontaneous EEG signals (δ, θ, α, β, and γ), while others may target evoked potentials. Long experiments may study changes in subjects' spontaneous EEG as they become more or less sleep-deprived, motion sick, etc. Evoked potentials have specific auditory or visual stimuli that target individual responses.

As mentioned earlier, EOG and EMG signals can also be recorded around the head, and provide a simple method of testing electrode connectivity.

Eye Blinks and EOG Artifacts

Spontaneous and large magnitude eye blinks can be recorded with almost any electrode on the scalp. Blink response amplitude often tracks with blink effort. Spontaneous and light blinks will cause minor deviations, while hard eye blinks can elicit large amplitude swings. All ear EEG measurements taken with WANDmini (including those used to verify the system in [35]) were initially tested with eye blinks, since they could help experiment hosts quickly determine adequate connectivity and acceptable noise/interference levels.Figure 5.3 shows eye blink amplitudes of about 0.2 mV to 1.2 mV (factor of 6) in the ear EEG case and about 1.8 mV to about 3.8 mV (factor of 2.1) when recorded with scalp EEG. The ear EEG recordings have an average baseline of 28 $\mu Vrms$ while the scalp EEG recording's average baseline is 16 $\mu Vrms$. Though the wet scalp electrodes recorded larger amplitude eye blinks (most likely due to the greater electrode spacing of scalp EEG), all grades of eye blinks are clearly visible when recorded with the ear EEG.

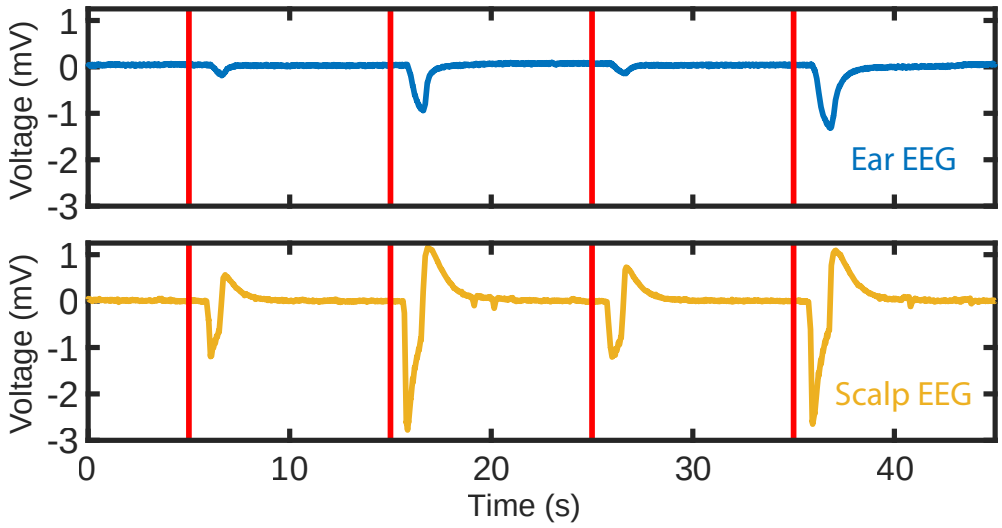

Figure 5.3: Eye blinks recorded from one subject with both the user-generic ear EEG and scalp EEG setups. Large blinks recorded with the user-generic ear EEG were 6x the amplitude of small blinks. Blinks recorded with the scalp EEG differ by 2.1x. Red lines mark visual cues.

Alpha Attenuation Response

Alpha rhythms are spontaneous neural signals between 8–13 Hz that reflect a person's state of attention or relaxation. Originating from the occipital lobe, alpha rhythms are commonly used to benchmark EEG systems due to their prevalence and large amplitudes. As a result, the alpha band modulation has been extensively researched as a possible BCI feature [14]. Moreover, alpha rhythms have become useful signals for drowsiness detection and other neurofeedback applications [82]. Alpha band power is directly modulated when a person closes their eyes. Furthermore, user-to-user variability in the alpha wave amplitude is well-documented [83]. To monitor the alpha wave with the ear EEG, subjects should sit comfortably in a quiet room and switch between two states at set intervals (e.g., every 30 s over 2 min): an eyes open/focused state and an eyes closed/relaxed state.

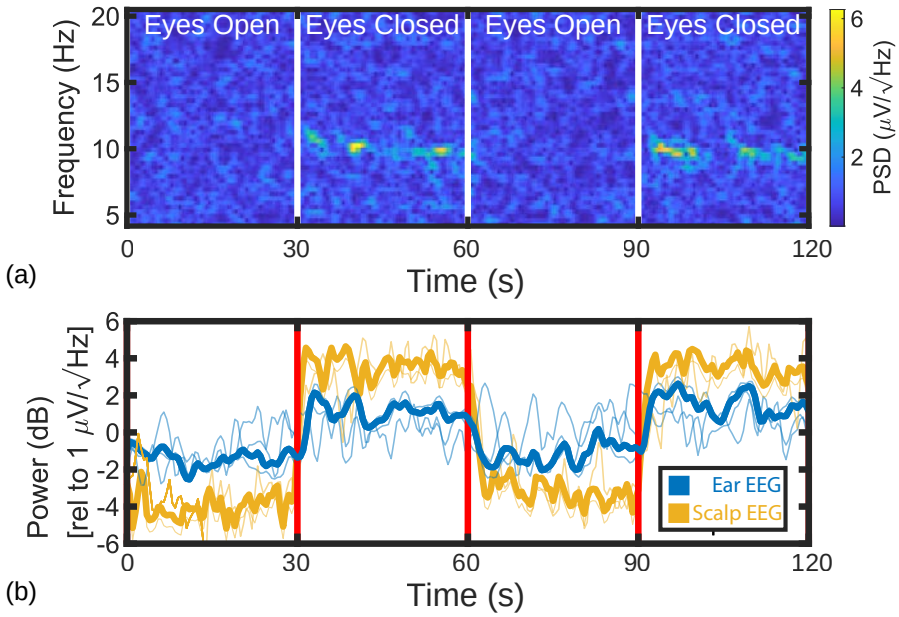

Figure 5.4: (a) Time-frequency spectrogram of alpha modulation recorded with the user-generic ear EEG on subject 2. Alpha (8–13 Hz) power increased by a factor of 3.4 in the eyes-closed state. (b) Individual averages (light traces) and grand average (bold trace) of mean alpha power for three subjects and 44 total trials recorded with user-generic ear EEG and scalp EEG.

When testing the ear EEG system from [35], the alpha attenuation response was the first neural signal robustly recorded. The spectrogram in Fig. 5.4(a) shows a representative example of a single user's alpha modulation recorded using the user-generic ear EEG. Fig. 5.4(b) shows the grand average of mean alpha (8–13 Hz) power from all three subjects. Mean alpha modulation R_{AM} is defined as the ratio of mean alpha power from eyes closed

to eyes open:

$$R_{AM} = \frac{P_{avg}\text{Alpha Band}_{\text{Eyes Closed}}}{P_{avg}\text{Alpha Band}_{\text{Eyes Open}}}. \tag{5.4}$$

The grand average alpha modulation (standard deviation) was 2.17 (± 0.69) (V^2/V^2) for all user-generic ear EEG subjects and 5.54 (± 1.84) (V^2/V^2) for wet scalp EEG. Specific subject and grand average alpha modulation ratios (with standard deviations) are shown in Table 5.1 and in Fig. 5.4(b). As expected, the scalp EEG recorded modulation is greater [20], but the signal maintains sufficient SNR for user-generic ear EEG detection.

Table 5.1: Individual and grand average alpha bandpower measurements

	Ear EEG	Scalp EEG
Subject 1	1.51 ± 0.17	5.29 ± 1.5
Subject 2	2.44 ± 0.62	5.71 ± 1.71
Subject 3	1.47 ± 0.54	3.06 ± 1.53
Average	2.17 ± 0.69	5.54 ± 1.84

Auditory Evoked Potentials

Perhaps the simplest evoked potentials to record are auditory evoked steady-state potentials (ASSR). Originating from the auditory cortex, ASSR can be evoked with different stimulus patterns, such as 40 Hz clicks, amplitude modulated white noise, or a 40 Hz amplitude-modulated 1 kHz tone [84] (Figure 5.5). In addition, recent work has shown that ASSR is a powerful tool for performing automatic hearing threshold estimation in infants [85] and determining binary choice via left-vs-right ear focus [86].

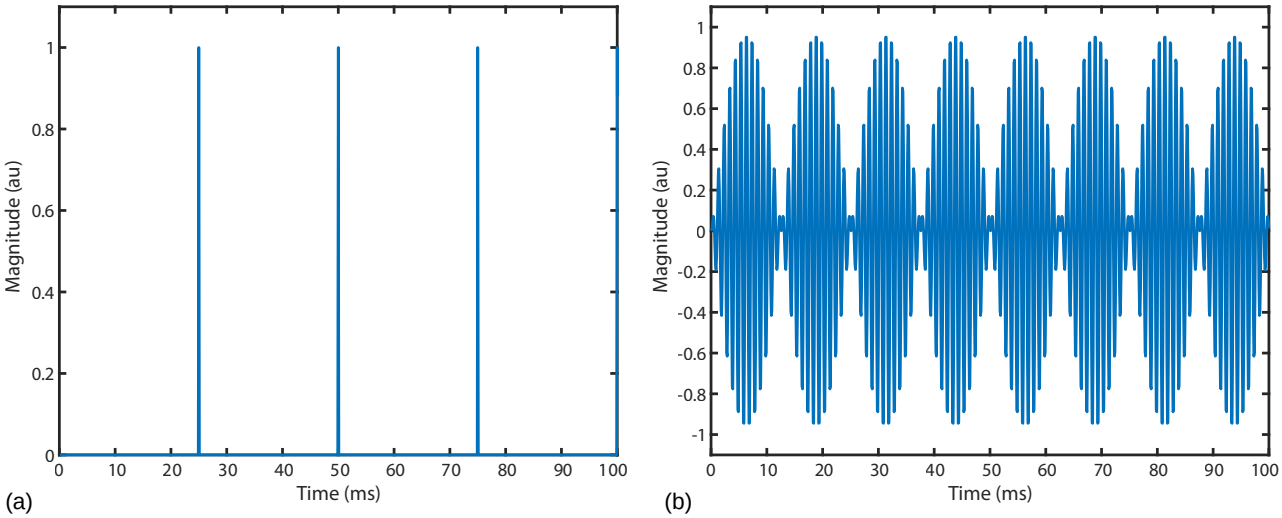

Figure 5.5: Example ASSR stimulus waveforms. (a) 40 Hz Clicks (b) 40 Hz amplitude modulated 1 kHz tone.

ASSR, a neural response centered at the same frequency as a low-frequency auditory stimulus, was targeted to confirm the ability to record evoked potentials with the user-generic ear EEG. To perform an ASSR experiment, each subject sat in a quiet room and listened to 40 Hz clicks sampled at 10 kHz played through a desktop speaker for 100 s. Mechanical coupling between the speaker and subject was minimized by placing a desktop speaker on a table in front of the user. This ensures any recorded ear EEG signal is indeed neural in origin. Fig. 5.6(a) and (b) show the grand averages of data recorded with the user-generic ear EEG and scalp EEG, respectively. All power spectral densities show a clear neural response at both 40 Hz and the second harmonic at 80 Hz without time-domain averaging. This lack of major post-processing and ease of stimulation makes ASSR another straightforward signal to the target. Furthermore, by virtue of being an auditory stimulus, the photosensitivity risk is minimized.

The mean SNR was 5.94 dB for user-generic ear EEG and 10.5 dB for scalp EEG. Note that inter-user variability exists between users due to differences in their hearing ability, thus,

the user-specific mean SNRs are shown in Table 5.2 (and are also represented in Fig. 5.6). Mean SNR was calculated as the ratio of the power at 40 Hz to the mean power from 35–45 Hz, excluding 40 Hz [20]:

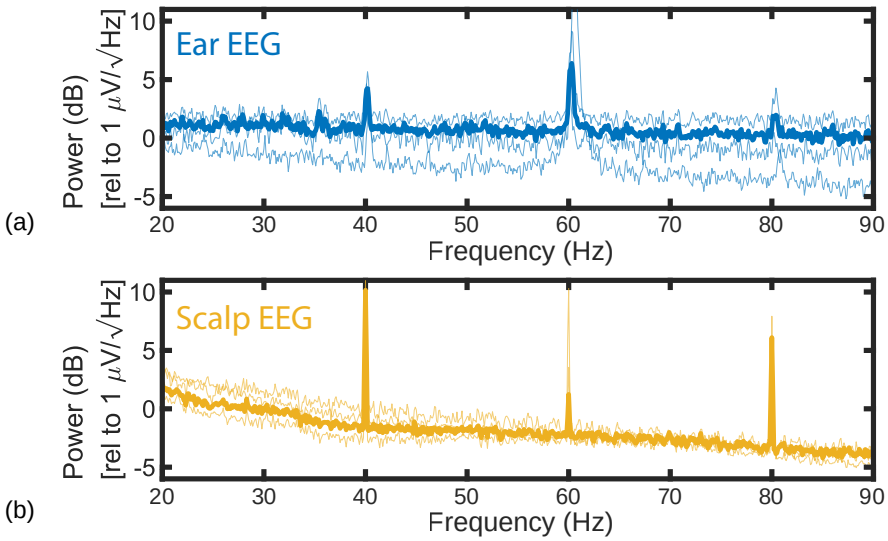

Figure 5.6: Auditory evoked potentials (ASSR for 40 Hz click stimulus) grand and individual averages. (a) Grand average PSD of ASSR for three subjects and 41 total trials for user-generic ear EEG showing a mean SNR of ~5.94 dB; (b) scalp EEG showing a mean SNR of ~10.5 dB.

Table 5.2: Individual and Grand Average ASSR 40 Hz SNR ± Std Dev (dB)

	Ear EEG	Scalp EEG
Subject 1	5.64 ± 1.39	13.79 ± 6.08
Subject 2	5.40 ± 0.99	7.36 ± 2.71
Subject 3	6.80 ± 2.02	8.51 ± 2.42
Average	5.94 ± 1.70	10.48 ± 5.23

While the SNR is lower for user-generic ear EEG (due to higher noise floor and shorter electrode distance, which reduces signal amplitude [20]), the SNR is sufficient for user-generic ear EEG detection. Furthermore, this ASSR signal was recorded across a single earpiece. Greater SNR could likely be achieved by recording across both ears.

Visually Evoked Potentials

Steady-state evoked potentials can also be evoked visually (visually evoked steady-state responses - VSSRs). Similar to ASSR, VSSRs can be elicited by showing a subject a screen flashing at a fixed rate between 10-30 Hz. It has been shown that if a subject has two monitors in front of them and each one is flashing at a different rate, then the recorded VSSR's spectral density will peak at the frequency that the subject is focusing on. While VSSR was not recorded with the Ear EEG system from [35], other in-ear EEG systems have verified its recordability [20]. It should be noted that VSSR experiments may cause discomfort for photosensitive subjects. Before running experiments, hosts should ask subjects if they are prone to seizures.

Steady State Response vs. Event Responses

Steady-state evoked potentials can only be viewed in the frequency domain after applying a Fourier transform. Event-related responses (ERPs), on the other hand, can be viewed in the time domain. All ERPs involve some state of recognition or surprise to sporadic stimuli. The most common experiment is referred to as an oddball paradigm where a subject's EEG is recorded when they are shown/hear different visual or auditory stimulus [26]. Often times there will be two classes of stimuli, an expected class and an oddball class. For example, the subject may be told to listen for multiple different tones, a 17 Hz beep or an 83 Hz beep (the frequencies are arbitrary). Every 20 seconds, either tone may play out of a speaker with the 17 Hz tone playing 80% of the time and the 83 Hz tone playing 20% of the time. If the 20-second epoch containing a 17Hz tone is averaged, there should be a specific p300 response, such as the example in Figure 5.7. In theory, a p300 should be a positive deviation from 0 roughly 300 ms after the stimulus onset. In practice, there is significant user-to-user variability, and the p300 can occur anywhere from 300 ms - 600 ms after the stimulus event.

In order to properly record ERPs, the stimulation signal and recording system must be perfectly synchronized such that the desired ERP's SNR increases when averaging. This synchronization can be difficult to ensure experimentally if the recording system does not allow for it. Oftentimes, recording systems have a 'synch' input or output such that they can trigger when stimulation signals occur or mark the succeeding samples.

The Power of Averaging

Signal averaging is a time domain processing technique used to increase the signal strength relative to any noise. Averaging a set of measurements taken with the same conditions can increase the recording's signal-to-noise (SNR) in proportion to the square root of the total number of averaged measurements. A derivation is below.

Assuming a per-sample noise variance of σ^2, averaging this random variable's variance n

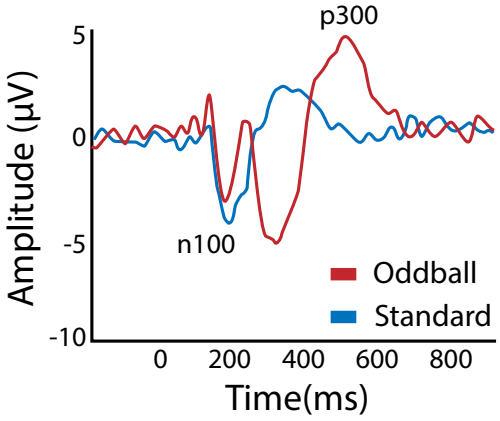

Figure 5.7: Example oddball paradigm with p300 response. The standard stimulus elicits a baseline response. Oddball stimuli exhibit a distinct response with a positive peak roughly 300 ms after the stimulus event.

times yields:

$$Var(\frac{1}{n}\sum_{i=1}^{n} z_i) \tag{5.5}$$

This can be simplified to:

$$\frac{1}{n^2}\sum_{i=1}^{n} Var(z_i) \tag{5.6}$$

Since the noise variance is a constant σ^2, the expression further simplifies to:

$$\frac{1}{n^2}\sum_{i=1}^{n} Var(z_i) = \frac{1}{n^2} n\sigma^2 = \frac{1}{n}\sigma^2 \tag{5.7}$$

This means noise power is reduced by a factor of n and that noise amplitude is reduced by a factor of $\sqrt{n}$.

Chapter 6

Hearable Application Case Study

Having demonstrated that ear EEG platforms can reliably record signals within the EEG band, the next step was developing an actual application. Ideally, the process of application development will also inform further system design optimizations as well. In this section, we will walk through the development of a pilot study targeting drowsiness detection and highlight how it can help us iterate on system design. This work was performed in collaboration with Carolyn Schwendeman [81]. All subject data was recorded in accordance with UC Berkeley's Institutional Review Board (CPHS protocol ID: 2018-09-11395). All volunteers were required to have normal vision and hearing abilities. All subjects were of 18-27 years of age.

While this report will not go over machine learning basics or any statistical methods, some system-level considerations will be reviewed. From an engineer's perspective, it is important to have a realistic expectation of how much training data can be recorded. In other words, the amount of training data will have large implications for which type of classifier model may be feasible to use. In addition, system design can benefit from prior knowledge on class separability (i.e. are classes linearly separable), common input features, and whether or not the classifier will be used as a future research tool. Neural network-based models, for example, generally require large training sets whereas linear regression-based models can be trained with small training sets. If there are known signal features that can separate between the desired classes, e.g. some statistical measure that perfectly separates class A from class B, then that will also impact model selection. Classical models such as support vector machines and random forest models generally work better with processed feature inputs as opposed to raw inputs, but neural networks can be incredibly effective with raw data. Lastly, if the system is meant to inform other applications or future systems, interpretable classifiers such as linear regression-based models should be prioritized. Interpretable models expose feature weights that can be directly analyzed to inform improvements in electrode mapping or electrical specifications.

6.1 Drowsiness Detection as an Application

Drowsiness has been linked to a decline in productivity, response time, and cognitive performance. When cognition is impaired, drivers, pilots, security guards, and personnel at nuclear facilities are all at increased risk of errors, occupational accidents, and bodily harm. In some cases, these accidents are fatal since sleepy and fatigued drivers account for up to 18% of vehicle accidents.

To prevent such accidents, drowsiness monitoring and detection have been implemented using camera-based eye tracking, steering trajectory sensors, and physiological recording devices. Camera-based systems rely on tracking the percentage of time the eyelids are closed, making them susceptible to obstructions. Steering trajectory sensors track steering wheel movement, limiting their use in cars and making them vulnerable to noise from rough roads. User-centered recording modalities including electrocardiography, electrooculography, and electroencephalography (EEG) are becoming increasingly popular since they can be taken with users and fit any work environment. Of these three bio-sensing paradigms, EEG features tend to achieve the highest accuracy in drowsiness detection.

Drowsiness detectors also share similar challenges to other longitudinal recording applications such as sleep, seizure, and focus monitoring. All of these applications require comfortable systems that can be worn repeatably, daily, and without the help of a technician. These electrodes should retain their electrical performance without serious servicing. Furthermore, the earpiece materials should not result in adverse reactions after use. In addition, drowsiness is significantly easier to induce and (in theory) monitor than epilepsy, sleep and focus. Monitoring drowsiness is thus need-driven, straightforward to monitor, and also informative for other applications such as epilepsy, sleep, and focus monitoring.

6.2 Experiment Design and Overview

The goal of each trial was to record a subject's EEG before, during, and after the onset of drowsiness. This was achieved by taking long EEG measurements while the user performed a tedious and uninteresting task in a comfortable environment. Each drowsiness trial was 40 – 50 minutes in length and took place indoors between 8 am and 5 pm with daytime lighting. Prior to each trial, users familiarized themselves with the experiment format and recording setup. Since behavioral and response-time measures are highly correlated to fatigue [9], task-based reaction times and user-reported Likert items were recorded during each trial and used to assign alert and drowsy labels during post-processing.

Users began each trial non-drowsy and seated in front of a laptop. Since the goal of each trial was to record drowsiness onset, users were instructed to play a repetitive game that measured their reaction time to cues on a graphical user interface (GUI). Every 60 seconds, a user was prompted to press a random number between 0-9 (Fig. 6.2a) and their reaction time was recorded. Every five minutes, the user was prompted to enter a Likert item according to the Karolinska Sleepiness Scale (KSS). This scale ranges from 0 = "extremely alert", to 10

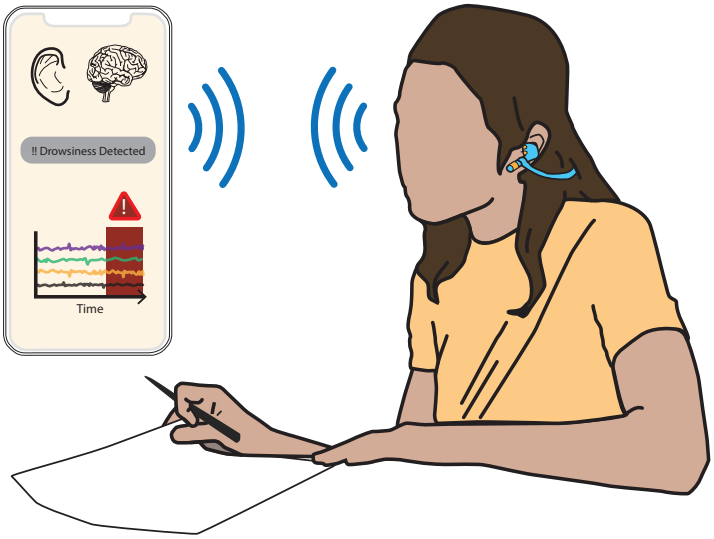

Figure 6.1: Envisioned Ear EEG device discreetly recording EEG from inside the ear canal and performing drowsiness detection for a user.

= "extremely sleepy, fighting to stay awake". EEG was recorded with Ear EEG throughout the trial. During initial tests, it became clear that recording from both ears to capture signals across the scalp improved the classifier accuracy by up to 10%. Thus, users wore two user-generic, dry electrode earpieces (one in each ear) and a compact wireless neural recording module WANDmini [35]. This experiment is based on [9] [87].

Twenty-one drowsiness trials were recorded with the Ear EEG system across five healthy users. To enable both user-generic and user-specific model training, users participated in multiple trials (maximum of five). Each trial exhibited at least one drowsy event lasting more than nine minutes. Thirty-six drowsiness events were recorded across all 21 recordings.

6.3 Data Labeling, Post-Processing, and Feature Extraction

After the trial, alert and drowsy labels were assigned based on reaction times and Likert items. Initial drowsy periods were highlighted anytime a user's reaction time exceeded 2.5 seconds. Then to account for momentary distractions where a user was alert but not looking at the GUI, any period that was not accompanied by a drowsy Likert item was relabeled as an alert period. Lastly, the starts and ends of drowsy periods were tuned based on the derivatives of a user's reaction times and Likert items. When the derivatives of the reaction time and Likert items increased, this marked the onset of a drowsy event. When

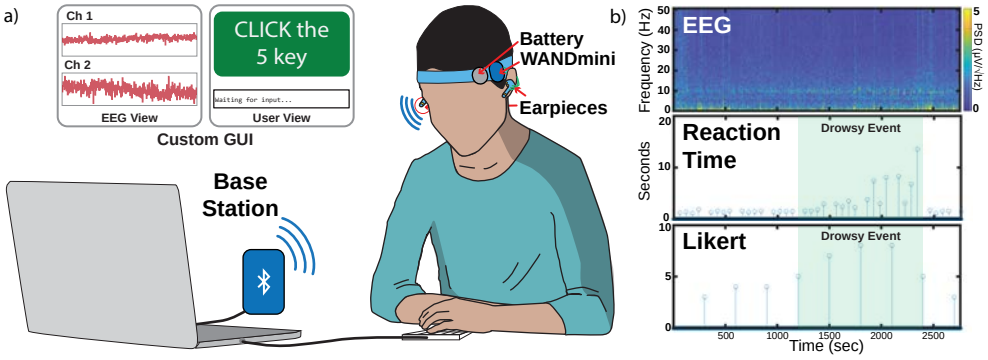

Figure 6.2: (a) Experimental setup. WANDmini records and transmits EEG from contralaterally worn earpieces to a base station via BLE. All EEG is live plotted while a GUI records the subject's reaction times and Likert items. (b) Sample spectrogram of single trial's EEG, reaction times, Likert items. Drowsy event in green.

the derivatives decreased, this marked the end of a drowsy event (Figure 6.2b). No EEG data was used to generate these labels.

Post-processing, feature extraction, and classification of the Ear EEG recordings were carried out to detect alert and drowsy states (Fig. 6.4). To take advantage of contralateral recordings, each in-ear electrode was re-referenced to the left concha cymba electrode to maximize spatial covering. Both originally referenced and re-referenced EEG channels were bandpass filtered from 0.05-50Hz in order to target the EEG frequency bands of interest for drowsiness detection. Namely, delta (δ, 0.05-4Hz), theta (θ, 4-8Hz), alpha (α, 8-13Hz), beta (β, 13-30Hz), and gamma (γ, 30-50Hz) [9].

Experimental recordings were divided into 50-second epochs, which began 10 seconds after a reaction time cue and ended when the next reaction time cue was provided. This focused classification on alert and drowsy states by removing EEG artifacts related to decision-making and motor planning in response to GUI cues. If a user's reaction time exceeded 10 seconds, the epoch was considered "sleep" and discarded from this alert and drowsy study. Features were calculated for each 50-second epoch since maximizing the feature extraction window size maximized classifier accuracy.

Both temporal and spectral features that appear in EEG-based classification literature were implemented [9] [87]. Time-domain features included voltage standard deviation and maximum peak-to-peak voltage amplitude. For frequency characteristics, the power spectral density (PSD) is calculated using Welch's method. Spectral features including maximum PSD, frequency of maximum PSD, and PSD variance were calculated for each EEG band. Absolute and relative band powers were also calculated for the following bands and ratios: δ, θ, α, β, γ, α/β, θ/β, $(\alpha+\theta)/\beta$, and $(\alpha+\theta)/(\alpha+\beta)$. The previous epoch features

were used to capture changes between subsequent epochs related to drowsiness onset. The relative bandpowers, bandpower ratios, and previous epoch features were found to be high-value observations for classifier training based on existing literature and also preliminary experiments. All three classes of features had some inherent normalization since they enable ratiometric classifications (i.e. they are resilient to cross channel/band interferers). Furthermore, using the previous epoch's features also enabled recognizing abrupt changes in feature values (i.e. alert to drowsy related changes). To account for outliers and signal artifacts, features were scaled by subtracting the median and scaling according to the interquartile range. During training, analysis of variance feature selection was used to find the 20 features that maximize class variation and minimize redundancy.

6.4 Classifiers

Lightweight, interpretable classifiers were selected due to the drowsiness dataset's small size. Furthermore, being able to sweep hyper-parameters such as channel count, feature count, and feature extraction window size, are paramount to understanding exactly how complicated the earpiece, readout, and classifier model need to be to achieve adequate performance. As a result, three models were selected to perform binary (alert/drowsy) classification: a logistic regression-based classifier, a support vector machine, and a random forest classifier. Information and background on these classification methods can be found in [88][89][90].

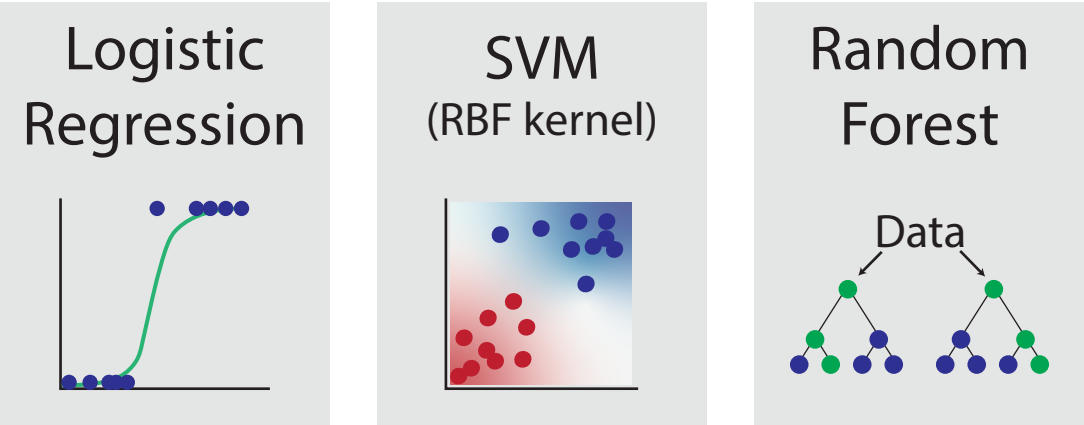

Figure 6.3: Basic diagrams of the three low-complexity classifiers selected to perform drowsiness detection.

All of these classifiers are trained individually and output a drowsiness probability for each feature window (50 seconds). Logistic regression is implemented with a stochastic average gradient descent solver and an L1 regularization, which adds a penalty equal to the

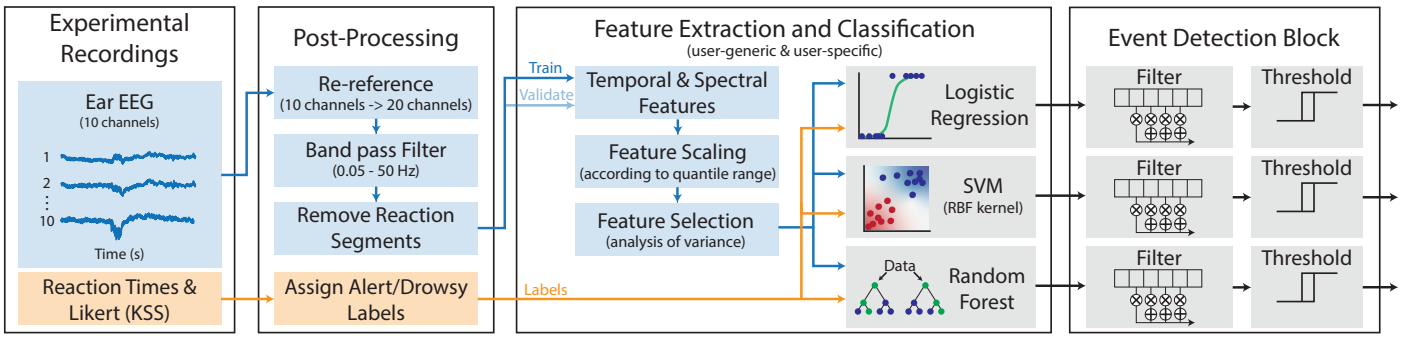

Figure 6.4: Block diagram outlining Ear EEG experimental recordings, post-processing, feature extraction, and classification to estimate alert and drowsy states. Reaction times and Likert items recorded during the trial are used to assign alert and drowsy labels for classifier training.

absolute value of the magnitude of feature coefficients. A support vector machine (SVM) is implemented with a radial basis function (RBF) kernel that utilizes approximately 170 support vectors and a regularization parameter, C=1. A random forest classifier is implemented with 100 trees and a maximum depth of five.

6.5 Event Detector

To reduce noise in the classifier output, class probabilities returned from all models are inputted to an event detection block consisting of a Hamming window FIR filter, thresholding stage, and an offline event detector. The moving average filter is implemented with an N-tap Hamming window (but can also be tuned to use different windows) FIR filter where the number of windows it filters (taps) can be tuned during training. This simple moving average filter greatly improves classification accuracy due to its elimination of single window alert/drowsy predictions. The output of this filter is then fed into a tunable thresholding stage (also set during training). The offline event detector blocks all of the contiguous non-drowsy and drowsy windows into larger blocks. Non-drowsy blocks are set to be 3 minutes in length (smaller blocks are ignored) while the drowsy event is chunked into a single block. Detection accuracy is then calculated based off if the average value of the block is consistent with that block's label. While this event detection strategy introduces a detection latency, it is significantly more robust to window-to-window noise. Furthermore, it is more realistic given that drowsiness is a slowly changing phenomenon.

6.6 Classification Results

Alert and drowsy event detection performance was estimated with three cross-validation techniques of varying practicality: user-specific, leave-one-trial-out, and leave-one-user-out. Commensurate with the state-of-the-art, each classifier's cross-validation performance was scored using event-based sensitivity (correctly classified drowsy events), specificity (correctly classified alert events), and accuracy (eq 1).

$$Accuracy = \frac{True\ Drowsy + True\ Alert}{All\ Events} \tag{6.1}$$

To consider a drowsy event correctly detected, classifiers must detect at least three minutes of consecutive 'drowsy' epochs. This threshold makes classification robust to noisy classifier outputs while ensuring acceptable detection latency. An average detection latency of 3.25 minutes was achieved across all training scenarios. Since this is before the user falls asleep, this latency is acceptable.

User-specific cross validation estimated event detection accuracy with recordings from individual users. The model was trained on n-1 trials for a user and tested on their remaining trial. This process was repeated n times and the results were averaged to estimate overall drowsiness detection accuracy. Average user-specific results ranged from 95.2 - 95.9% across all classifier models (Fig. 6.5). Leave-one-trial-out cross validation estimated event detection accuracy with recordings from all users. The model was trained on 20 trials and tested on the remaining trial. This process was repeated 21 times. A 93.9 – 95.4% average event detection accuracy was achieved across models. Leave-one-user-out cross validation trained the model on recordings from four users and tested on recordings from the never-before-seen user. This user-generic process was repeated five times. Across all classifier models, a 94.5 - 95.0% average event detection accuracy was achieved. These results imply that a monolithic user-generic model can be trained on a test population and then used in a deploy-and-forget drowsiness detector. This certainly would be the ideal case, but this is an unlikely outcome as the sample population grows. As will be discussed in later sections, even slightly larger population sizes will demonstrate degradation in accuracy when using shared blocks that are not tuned to maximize individual accuracy.

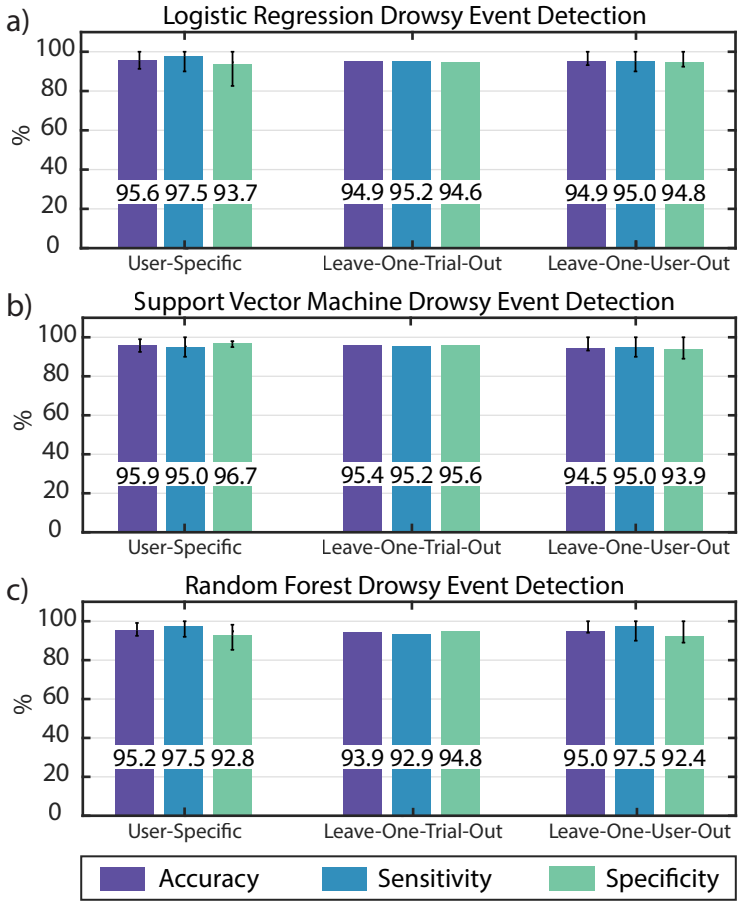

Figure 6.5: Drowsy-event detection across three classifier models (a) logistic regression, (b) support vector machine, (c) random forest. For each model, results are shown for user-specific, leave-one-trial-out, and leave-one-user-out cross-validation. Error bars show minimum and maximum user results.

6.7 Accuracy Impact Across Larger Test Populations

The monolithic model presented earlier works well when training and testing on a small population of five subjects, this is not guaranteed to remain the case as more subjects are added. It is likely that as the test population increases, more subjects will demonstrate different drowsiness 'signatures' (specific feature values that indicate specific brain states). For example, while all users may be able to utilize a single set of feature inputs, the event detector may need to be tuned for specific users since they may have shorter/different responses. In fact, this is the case as the subject pool increases from 5 users to 9.

To demonstrate the need for some tuning even in deploy-and-forget scenarios, the user-generic pipeline (leave-one-user-out training scenario) can be investigated across different training populations. A user-generic pipeline with shared event-detector settings would ideally require no tuning when testing on different users. With a 5 subject training population, this shared pipeline classifies with the following accuracies for 3 specific subjects (Table 6.1). To simplify comparisons, only the SVM performance will be investigated. With a shared event detector, all users achieve high sensitivity, specificity, and overall accuracy.

Table 6.1: 5 subject case: specific subject SVM performance with shared event detector

Case 1	Accuracy	Sensitivity	Specificity
Subject 1	95.8%	100%	95.4%
Subject 2	90.5%	100%	89.6%
Subject 3	85.0%	80.0%	86.6%

If we increase the subject pool to 9 subjects, retrain the classifier, and retune the event detector to maximize average accuracy, these specific subjects exhibit different performance than before (Table 6.2). Subject 1 exhibits the same performance, Subject 2 exhibits degraded performance, and Subject 3 exhibits better accuracy than before but with significantly degraded sensitivity. This is likely due to the fact that a shared event detector is unable to

Table 6.2: 9 subject case: specific subject SVM performance with shared event detector

Case 2	Accuracy	Sensitivity	Specificity
Subject 1	95.8%	100%	95.4%
Subject 2	88.5%	80%	89.3%
Subject 3	88.0%	70.0%	92.8%

handle the different classifier outputs for different subjects. Some subjects may have very well-defined drowsiness states where the classifier can detect the onset of drowsiness with high confidence. Other subjects may have more noise in their feature values which may result in lower confidence classifier outputs. If the event detector's thresholding stage is weighted more towards high confidence outputs, the noisier subject data (from poor electrode fit or differences in EEG responses) will perform significantly worse.

This theory is confirmed by tuning the event detectors to optimize for individual user performance as opposed to average user performance (Table 6.3). 6.3). By performing some

Table 6.3: 9 subject case: specific subject SVM performance with customized event detector

Case 3	Accuracy	Sensitivity	Specificity
Subject 1	95.8%	100%	95.4%
Subject 2	94.1%	100%	93.6%
Subject 3	84.5%	80.0%	86.0%

basic tuning of the event detector's threshold and averaging kernel, performance can be boosted across all three subjects. Subject 1 still performs well, Subject 2 performs better across all metrics than in previous cases, and subject 3 exhibits a 10% boost in sensitivity for a small hit in specificity. As a result, the earlier conclusion is that a deploy-and-forget system is possible across large test populations. While the majority of the classification pipeline can be shared for each user, it is important to have some user-specific tuning in the pipeline aside from training feature weights.

6.8 Hyperparamater Sweeps Can Inform Future Design Improvements

Small models can still provide insight into future system improvements and optimizations. Observing how model performance changes across different model hyperparameter sweeps (channel count, feature count, window size, etc.), can illuminate optimization opportunities in the earpiece design, electrode placement, and front-end requirements. Thus to inform further system iterations, sweeps were performed across the original five subject dataset. As shown before, the classifiers tend to perform similarly and a trend in one model is often found in the others. In other words, if a certain channel has greater SNR, it is often selected regardless of the model. As a result, the sweeps shown below are all with the SVM classifier to make the results and conclusions more interpretable.

A primary question is whether or not two earpieces are needed for drowsiness detection or if a single earpiece would suffice. The appreciable difference between a single earpiece system

and a two earpiece system is that the prior relies on sense-reference electrode pairs in the same ear, i.e. an ipsilateral measurement. While this was shown viable in [33] and [35], the SNR of ipsilateral measurements tends to be worse than for measurements taken across both ears, i.e. contralateral measurements. This is because the signal of interest will ideally have a significantly greater amplitude relative to a far away reference (contralateral) than a nearby reference (ipsilateral) [20]. This expectation is confirmed when the classifier only operates on channels from a single earpiece (Figure 6.6). While further optimization may help to

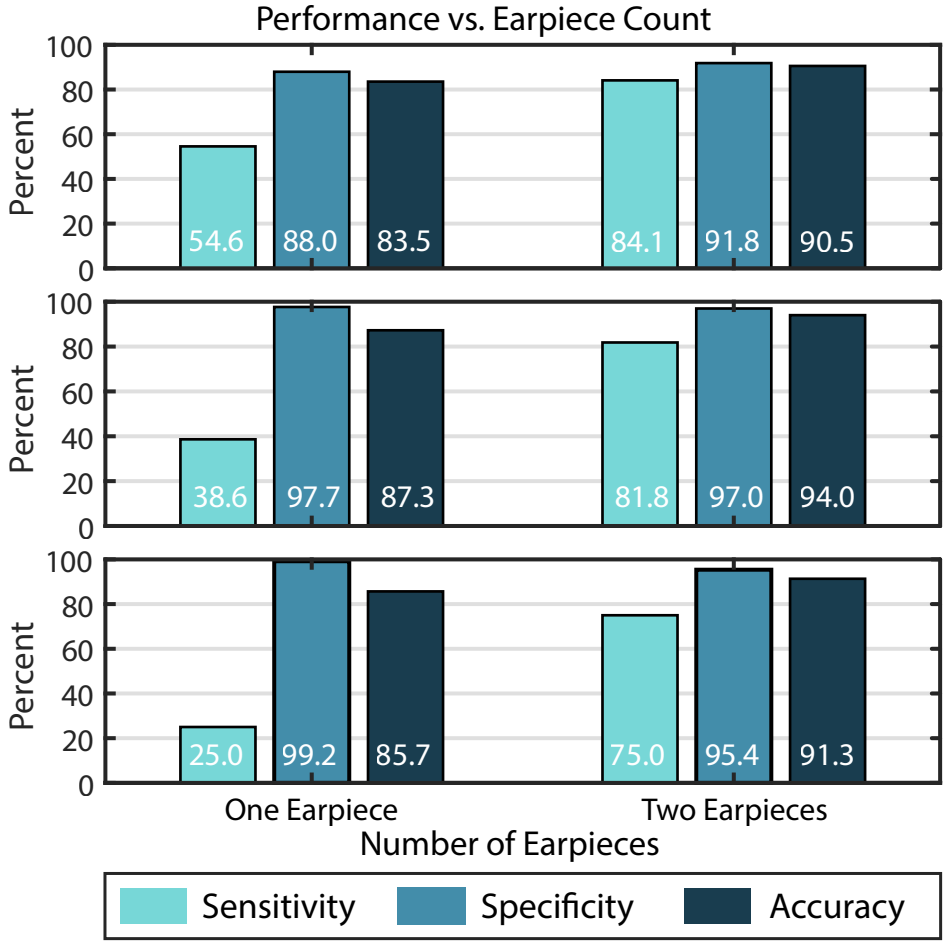

Figure 6.6: Average SVM-based classifier performance across single earpiece and two earpiece scenarios for a 5 subject population. Contralateral measurements taken with two earpieces greatly improve classifier performance across all training scenarios.

improve the accuracy of both cases, such improvements would not be able to overcome the 30 - 50% difference in accuracy between a single earpiece and two earpiece systems. This performance gap is consistent across all training scenarios and classifier models as well. As

a result, a future system should leverage contralateral measurements via two earpieces to maximize classifier performance.

It is important to note that a single earpiece system only has 10 channels maximum. Five physical sense-reference pairs and another five re-referenced, 'virtual' channels. As a result, it is imperative to confirm that the difference shown in Figure 6.6, is not related to the channel and feature count. This can quickly be demonstrated by sweeping the number of classifier input channels from 1 to 20 (Figure 6.7). One channel is a single sense-reference pair, where 20 channels include all physical channels (10) and their re-references to the left cymba electrode (another 10). In order to sweep the number of channels, the classifier pipeline first performs feature extraction and simple regression to estimate which channels result in the greatest output class variance. In other words, the feature extraction step highlights which channels have the greatest amount of unique information. An interesting phenomenon is

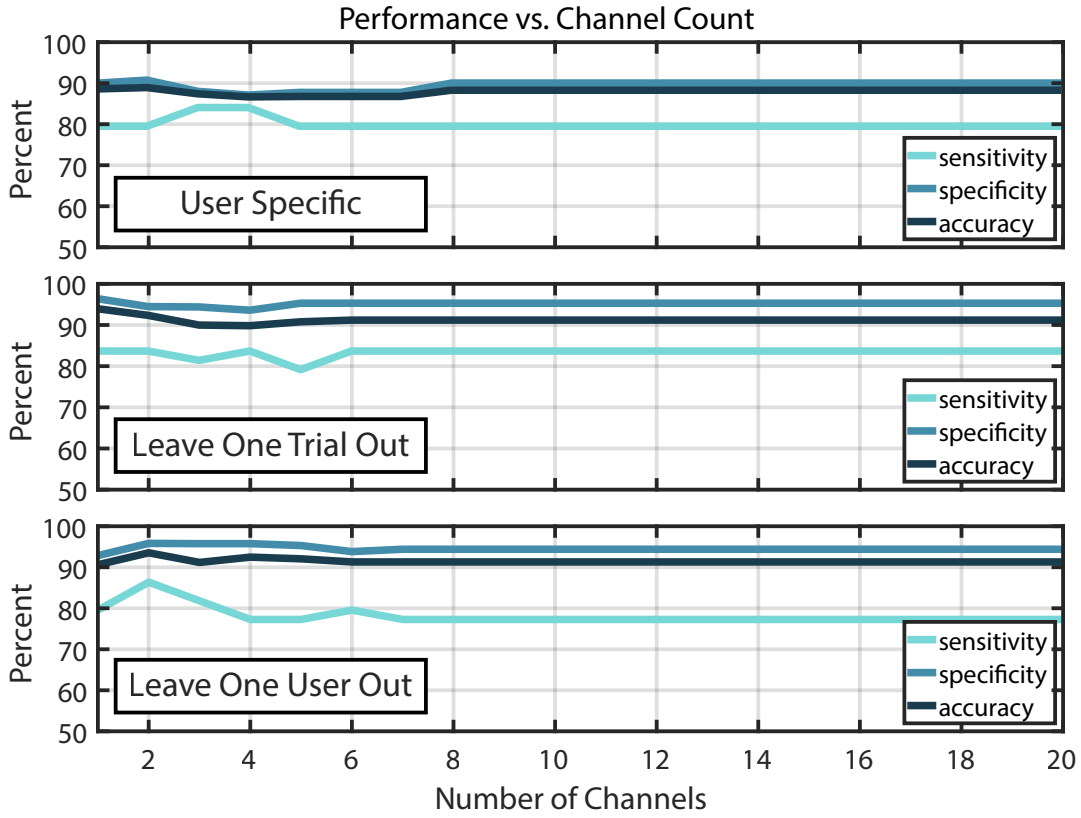

Figure 6.7: SVM-based classifier performance as the number of input channels are swept for a 5 subject population. Past 2-4 channels, there is little to no performance benefit across all training scenarios. This is likely due to the correlated nature of channels in close proximity to each other.

immediately clear, significantly fewer channels are required than the system specifications

suggest. While the Ear EEG system can support up to 10 physical channels, only one to three seem to be required for maximum accuracy. Furthermore, when seeing which channels are continuously weighted as the most important during the channel selection, contralateral channels are always in the top five. Not only does this information confirm the single vs. two earpiece conclusions, but it also highlights a clear optimization going forward. Fewer channels and electrodes can be used to achieve the best possible performance. There is one last point, it seems that increasing the channel count can sometimes actually hurt classifier performance, this is possible in cases where input observations/features are highly correlated to each other. This implies that most of the Ear EEG channels are highly correlated, which actually confirms the initial analysis performed by [35] and [20].

Related to channel selection, feature selection and count also have a large part to play in a system's power consumption. If a large number of features must be calculated for every classification, this implies more memory and computational blocks are required. Increasing the memory and computational requirements in turn also increases power consumption and system complexity. Figure 6.8 illustrates how accuracy changes as the total number of features inputted to the classifier models are increased from one to 10. This analysis shares some overlap with the channel count sweep due to the fact that it is now clear that fewer channels perform just as well as large channel count systems. As a result, this sweep can focus on miniaturizing the classifier pipeline further. Prior to each training and testing stage, this sweep performed feature selection where the top features (regardless of channel) that maximized class variance were selected. What resulted is consistent with the channel count sweep. Significantly fewer features than expected are required to achieve maximum accuracy. On top of that, the top selected features are all contralateral relative bandpower calculations where the delta and alpha bandpowers are divided by the total bandpower from 0-50Hz (this effectively normalizes the absolute bandpower calculations).

Figure 6.8 illustrates classifier performance across 1-10 input features to an SVM based classifier. It is important to note that the demonstrated behavior is consistent with even larger input feature counts. Figure 6.9 demonstrates classifier performance across 10-50 input features.

Another important system specification is the frequency at which classifications are performed. In other words, how large of a data window can be used for feature extraction and classification? Smaller data windows (e.g. 10 seconds), may use less memory than larger data windows (50 seconds) but will result in 5 times as many feature calculations and classifications as the latter. While the power considerations will be highly implementation specific, it is still important to understand the maximum or minimum window size because it would directly influence any detection latency as well. Testing this parameter is more straightforward than other sweeps in that all that is needed is a change to the window size and increment. Figure 6.10 shows sensitivity, specificity, and accuracy across different window sizes. There is no overlap between subsequent windows in the classification pipeline. While the user-specific training scenario seems to have higher sensitivity at smaller window sizes, the leave-one-user-out and leave-one-trial-out scenarios perform better with larger window sizes. Due to experimental restraints, 50 seconds was the largest possible window size. Other

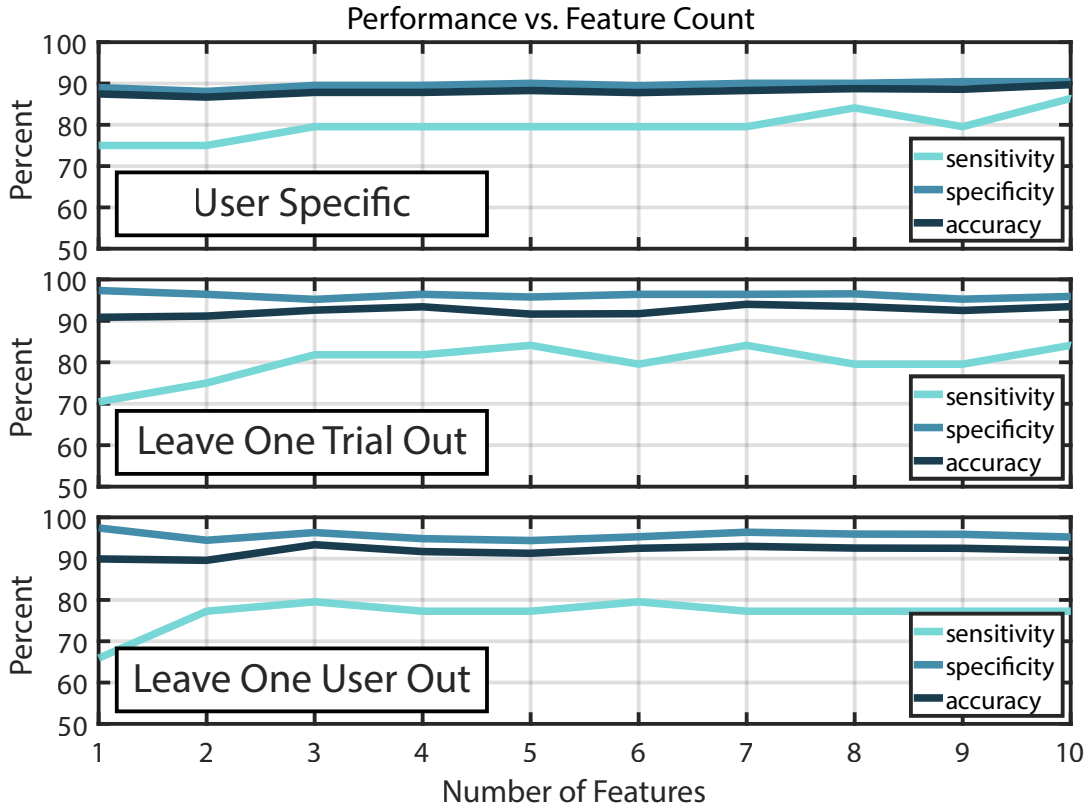

Figure 6.8: SVM-based classifier performance as the number of input features are swept from 1 to 10 for a 5 subject population. After 4-5 features, accuracy reaches an effective steady state across all training scenarios.

works such as [87] demonstrated that accuracy improves with window sizes greater than 200 seconds, this is likely due to drowsiness being highly correlated to low-frequency EEG signals. These low-frequency EEG signals are intermittent during drowsiness onset but can range from 1 Hz to 10 Hz. As a result, large window sizes improve the likelihood that clean, high SNR low-frequency features will be visible in the case of drowsiness specifically.

This analysis, even in its limited capacity, suggests that a far simpler system with fewer channels, fewer features, and a slower classification output rate could potentially achieve the same sensitivity, specificity, and accuracy as the presented Ear EEG system built with the user-generic earpiece and WANDmini. A new system design would focus on a new earpiece with fewer electrodes, a new readout with fewer channels, and a classifier design that leverages the slow nature of drowsiness features.

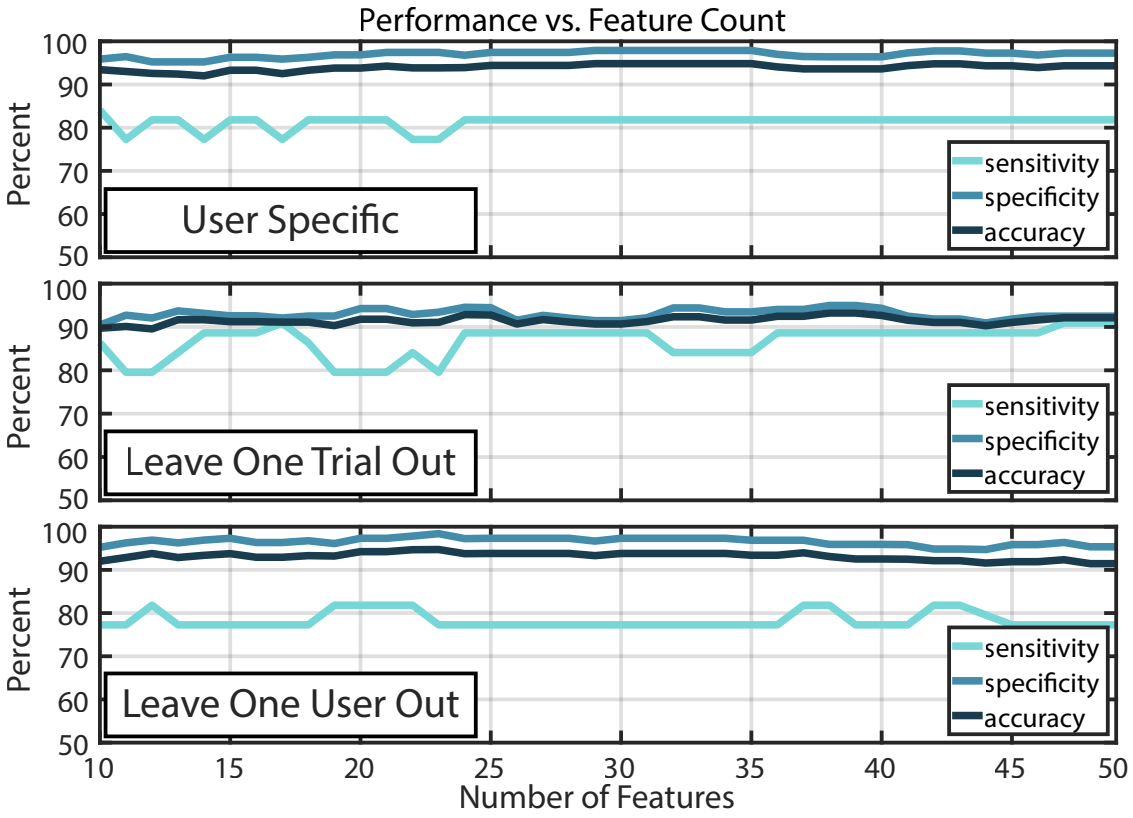

Figure 6.9: SVM-based classifier performance as the number of input features are swept from 10 to 50 for a 5 subject population. Confirming what was demonstrated in Figure 6.8, there is little performance benefit for increased feature counts. In the Leave One Trial Out case, an input feature count of around 25 seems to stabilize sensitivity.

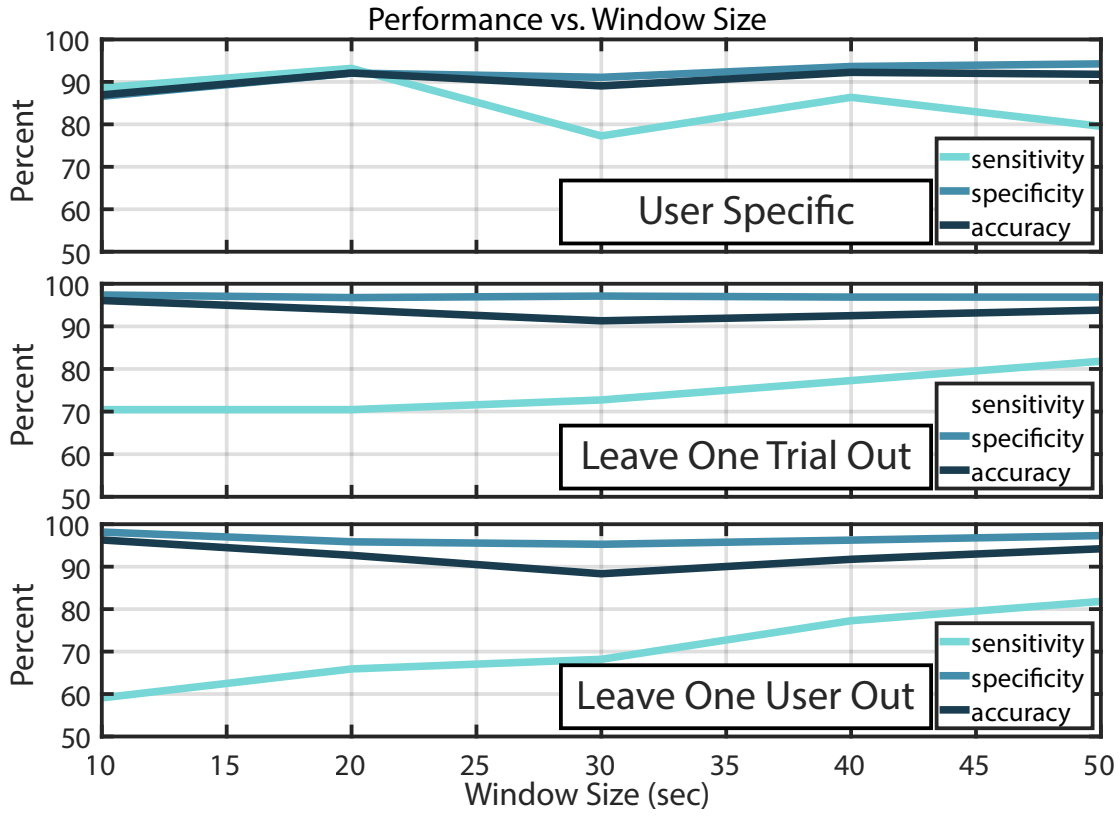

Figure 6.10: SVM-based classifier performance as the feature extraction window size is swept from 10-second windows to 50-second windows for a 5-subject population. While the user-specific case may benefit from smaller window sizes, the leave-one-trial-out, and leave-one-user-out cases to see 10 to 20% improvements in sensitivity.

Chapter 7

Summary: Now Can We Build Neural Wearables?

7.1 Overcome Barriers

Through this report, major parts of an Ear EEG Hearable have been discussed. Earpiece design and fabrication have been demonstrated with publically available anatomical databases and commercial fabrication processes such as 3D printing and electroless plating. The example earpiece and electrodes achieved relatively low electrode skin impedance across multiple users without skin preparation or cleaning. Commercial components and PCB services can produce wireless recording front ends and ambulatory, low-profile systems. The wireless readout, WANDmini, demonstrated long-term ($>$ 2 hour) experiments. Open source software tools such as scikit in python can enable machine learning-based applications. With the previously mentioned examples, a small batch of devices can be constructed to achieve high accuracy brain-state classification across multiple users over the course of multiple hours. Ultimately, an end-to-end Ear EEG device has been demonstrated, iterated on, and verified with multiple small population human subject trials.

7.2 Remaining Challenges and Possible Experiments

This work has established a baseline device that can be used in multiple wireless, ambulatory contexts across multiple different users and time periods. Furthermore, existing experiments have already proven to be fertile ground for improving system reliability, comfort, and robustness. Based on conclusions in the previous chapter, it is clear that an even simpler user-generic Ear EEG device with only 3-4 channels (Figure 7.1a) could record actionable data for a brain-state classifier. This simpler earpiece could use even larger flexible electrodes (potentially even flex PCBs on a silicone substrate) to further lower ESI and improve the probability of good in-ear electrode contact across more users (Figure 7.1b). This is because larger electrodes increase the chance of stable contact and greater contact areas lower ESI.

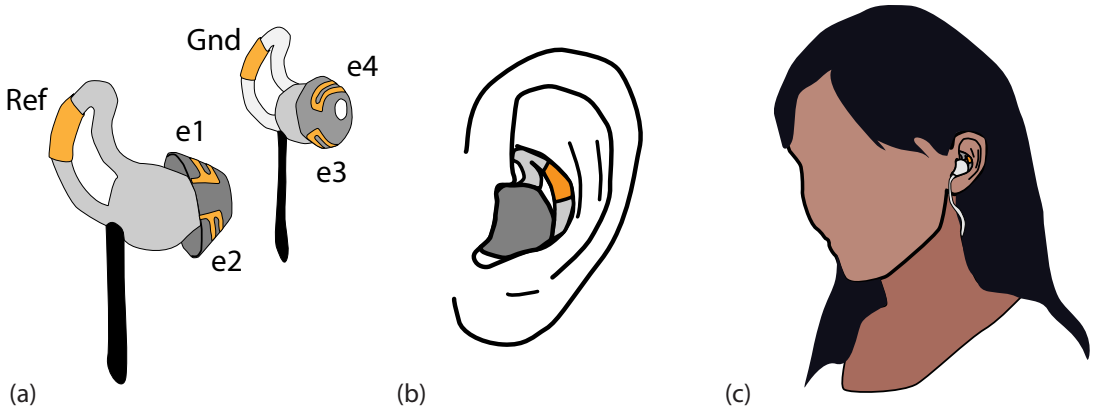

Figure 7.1: Envisioned future (a) earpieces and (b) integration that could enable (c) truly discreet neural recording through Ear EEG.

After constructing larger quantities of these simpler earpieces, the existing hardware could be improved to implement lower power operation (by reducing channel count) and adding features such as simultaneous impedance monitoring [71] and the latest wireless standards. By incorporating simultaneous impedance monitoring, Ear EEG systems could monitor electrode contact to grade neural recording quality and even possibly remove motion artifacts without reducing EEG SNR. Adding other features, such as improved Bluetooth stacks, can also help reduce power consumption and incorporate wireless security features to begin alleviating future privacy concerns. Ultimately, this optimized hardware could also be minimized to potentially fit inside the earpieces to achieve a truly discreet, integrated neural wearable (Figure 7.1c).

With new earpieces and hardware, another thrust would be to vastly increase the number of human subject trials to answer the growing list of questions regarding dry electrode Ear EEG use in daily life. What motion artifacts occur while driving cars, riding buses, or taking trains? Can motion artifacts be used to classify a user's actions such as eating or talking? Can Ear EEG devices use motion and EMG activity to detect the type of food someone is eating? Can Ear EEG devices be used to monitor EEG activity in the presence of any of the aforementioned activities? Can Ear EEG devices robustly record sleep data across multiple users? What other neural signals can be robustly recorded with Ear EEG across multiple users? All of these and many more questions require large amounts of control data which is now possible to begin recording. Many of these can be performed without extensive internal review board (IRB) experimental guidance, but IRB and medical guidance will certainly be required for many other experiments. Ideally, clinical studies could also be performed to explore Ear EEG's viability in long-term sleep and epilepsy monitoring.

Before commercialization and mass production of Ear EEG devices, a rigorous ethical

discussion should also be had. Never before has it been possible that the average user's brain signals could be recorded and digitized throughout their daily lives. While it currently seems that the SNR of recorded data limits the scope of Ear EEG applications to large amplitude features specifically related to sleep and epilepsy, future advances may unlock signals that potentially invade user privacy. As a result, it is important to take care of any recorded data, bolster the security of any wireless transfers, and be upfront with users specifically about what data is being recorded and how such measurements will be used.

It is important to note that many of these remaining problems are primarily engineering problems that can and should make use of existing solutions in other fields. At this point, it is unlikely new conductive materials are required and existing recording hardware techniques are already generally adequate. The largest remaining research questions are computational in nature (motion artifact removal with simultaneous ESI) and physiological (performing larger-scale human subject trials). When it comes to physically build these devices, most steps are solved problems with the proper expertise.

7.3 Parting Remarks

It is my belief that these devices, once made smaller and more comfortable, have the capacity to provide life-changing utility to some users. If seizures can accurately be detected or even predicted with ear EEG devices or users can augment their sleep patterns to improve their quality of life, then this is all incredibly worth it. What I've laid out, hopefully somewhat clearly, is a pathway to make these systems with relative ease using low-cost rapid fabrication techniques that are scalable. The example Ear EEG developed here shows that user-generic, wireless hearables really are possible! They can discreetly record neural activity with enough SNR to provide brain-state classification and fit an intuitive form factor. Dry electrodes remove any technician-related barriers to donning and doffing these wearables while utilizing earbud form factors means users will (ideally) know exactly how to wear such devices. At most, some simple pictures will remove any confusion on proper wear. The next steps are up to future students and entrepreneurs to scale it all up. Then maybe such Ear EEG devices will be used not just as drowsiness, sleep, and relaxation monitors, but also as truly life-saving devices in the form of seizure predictors.

Bibliography

[1] Jessilyn Dunn, Ryan Runge, and Michael Snyder. "Wearables and the medical revolution". In: *Personalized Medicine* 15 (5 Sept. 2018), pp. 429–448.

[2] D. J. McFarland and J. R. Wolpaw. "EEG-based brain–computer interfaces". In: *Current Opinion in Biomedical Engineering* 4 (Dec. 2017), pp. 194–200.

[3] Soheyl Noachtar and Jan Rémi. "The role of EEG in epilepsy: A critical review". In: *Epilepsy and Behavior* 15 (1 May 2009), pp. 22–33.

[4] Luz Maria Sanchez-Reyes et al. "Impact of EEG Parameters Detecting Dementia Diseases: A Systematic Review". In: *IEEE Access* 9 (2021), pp. 78060–78074.

[5] Alice Jackson and Donald J. Bolger. "The neurophysiological bases of EEG and EEG measurement: A review for the rest of us". In: *Psychophysiology* 51 (2014), pp. 1061–1071.

[6] Eric R. Kandel, James H. Schwartz, and Thomas M. Jessell, eds. *Principles of Neural Science*. Third. New York: Elsevier, 1991.

[7] Piotr Olejniczak. "Neurophysiologic Basis of EEG". In: *Journal of Clinical Neurophysiology* 23 (3 2006), pp. 186–189.

[8] Michael X. Cohen. "Where Does EEG Come From and What Does It Mean?" In: *Trends in Neurosciences* 40 (4 Apr. 2017), pp. 208–218.

[9] Shaibal Barua et al. "Automatic driver sleepiness detection using EEG, EOG and contextual information". In: *Expert Systems with Applications* 115 (Jan. 2019), pp. 121–135. URL: `https://linkinghub.elsevier.com/retrieve/pii/S0957417418304792`.

[10] Vojkan Mihajlovic et al. "Wearable, wireless EEG solutions in daily life applications: What are we missing?" In: *IEEE Journal of Biomedical and Health Informatics* 19 (1 2015), pp. 6–21.

[11] Terence W. Picton et al. "Estimating audiometric thresholds using auditory steady-state responses". In: *Journal of the American Academy of Audiology* 16 (3 2005), pp. 140–156.

[12] Xiyuan Hou et al. "EEG Based Stress Monitoring". In: *Proceedings - 2015 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2015* (2016), pp. 3110–3115.

[13] Gernot R. Müller-Putz and Gert Pfurtscheller. "Control of an electrical prosthesis with an SSVEP-based BCI". In: *IEEE Transactions on Biomedical Engineering* 55 (1 Jan. 2008), pp. 361–364.

[14] Gerwin Schalk et al. "BCI2000: A general-purpose brain-computer interface (BCI) system". In: *IEEE Transactions on Biomedical Engineering* 51 (6 June 2004), pp. 1034–1043.

[15] Matthieu Duvinage et al. "A P300-based quantitative comparison between the emotiv epoc headset and a medical EEG device". In: 2012, pp. 37–42.

[16] A. Searle and L. Kirkup. "A direct comparison of wet, dry and insulating bioelectric recording electrodes". In: *Physiol. Meas.* 21.2 (2000), pp. 271–283.

[17] Mullen Joellan and Wendy Morton. "Preventing Skin Breakdown in EEG Patients: Best Practice Techniques". In: *Journal of Pediatric Nursing* 29.5 (2014), pp. 478–480. ISSN: 0882-5963.

[18] M. Teplan. "Fundamentals of EEG measurement". In: *Meas. Sci. Rev.* 2 (2002), pp. 1–11.

[19] William O Tatum, Barbara A Dworetzky, and Donald L Schomer. *Artifact and Recording Concepts in EEG*. 2011, pp. 252–263.

[20] Simon L. Kappel et al. "Dry-Contact Electrode Ear-EEG". In: *IEEE Trans. Biomed. Eng.* 66.1 (2019).

[21] M. A. Lopez-Gordo, D. Sanchez Morillo, and F. Pelayo Valle. "Dry EEG Electrodes". In: *Sensors* 14.7 (2014), pp. 12847–12870.

[22] Federico Tasca et al. "Microneedle-based electrochemical devices for transdermal biosensing: a review". In: *Current Opinion in Electrochem.* 16 (2019), pp. 42–49. ISSN: 2451-9103.

[23] Jiahui Xu and Baichang Zhong. "Review on portable EEG technology in educational research". In: *Computers in Human Behavior* 81 (2018), pp. 340–349. ISSN: 0747-5632.

[24] D. Looney et al. "An in-the-ear platform for recording electroencephalogram". In: *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS* (2011), pp. 6882–6885.

[25] Preben Kidmose et al. "A study of evoked potentials from ear-EEG". In: *IEEE Transactions on Biomedical Engineering* 60 (10 2013), pp. 2824–2830.

[26] Arnd Meiser and Martin G. Bleichner. "Ear-EEG compares well to cap-EEG in recording auditory ERPs: A quantification of signal loss". In: *Journal of Neural Engineering* 19 (2 Apr. 2022).

[27] "Ear-EEG detects ictal and interictal abnormalities in focal and generalized epilepsy – A comparison with scalp EEG monitoring". In: *Clinical Neurophysiology* 128 (2017), pp. 2454–2461.

[28] I. C. Zibrandtsen, P. Kidmose, and T. W. Kjaer. "Detection of generalized tonic-clonic seizures from ear-EEG based on EMG analysis". In: *Seizure* 59 (2018), pp. 54–59.

[29] Takashi Nakamura, Valentin Goverdovsky, and Mary J Morrell. "Automatic sleep monitoring using ear-EEG". In: 5 (March 2017).

[30] David Looney et al. "The in-the-ear recording concept: User-centered and wearable brain monitoring". In: *IEEE Pulse* 3 (6 2012), pp. 32–42.

[31] Kaare B. Mikkelsen et al. "EEG recorded from the ear: Characterizing the Ear-EEG Method". In: *Frontiers in Neuroscience* 9 (Nov 2015).

[32] Valentin Goverdovsky et al. "In-Ear EEG From Viscoelastic Generic Earpieces: Robust and Unobtrusive 24/7 Monitoring". In: *IEEE Sensors Journal* 16 (1 2016), pp. 271–277.

[33] Ryan Kaveh et al. "A Wireless, Multielectrode, User-generic Ear EEG Recording System". In: *2019 IEEE Biomedical Circuits and Systems Conference (BioCAS)* (2019). URL: `https://ieeexplore-ieee-org.libproxy.berkeley.edu/document/8918700`.

[34] Nicole Gleichmann. *Hearing Aid Prices: How Much Do Hearing Aids Cost?* 2022. URL: `https://aginginplace.org/hearing-aids/prices/`.

[35] Ryan Kaveh et al. "Wireless User-Generic Ear EEG". In: *IEEE Transactions on Biomedical Circuits and Systems* 14.4 (Aug. 2020), pp. 727–737.

[36] "A new soft material based in-the-ear EEG recording technique". In: *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS* 2016-October (2016), pp. 5709–5712.

[37] Akshay Paul et al. "Attention State Classification with In-Ear EEG". In: Institute of Electrical and Electronics Engineers Inc., 2021.

[38] A R Bertelsen et al. *Generic Dry-Contact Ear-EEG; Generic Dry-Contact Ear-EEG*. 2019.

[39] Allen J. Bard and Larry R Faulkner. *Electrochemical Methods : Fundamentals and Applictions*. 1980.

[40] David A. Stenger and Thomas M. McKenna. *Enabling Technologies for Cultured Neural Networks*. 1994.

[41] "Electron transfer processes occurring on platinum neural stimulating electrodes: Calculated charge-storage capacities are inaccessible during applied stimulation". In: *Journal of Neural Engineering* 14 (4 May 2017).

[42] Daniel R. Merrill, Marom Bikson, and John G.R. Jefferys. "Electrical stimulation of excitable tissue: Design of efficacious and safe protocols". In: *Journal of Neuroscience Methods* 141 (2 Feb. 2005), pp. 171–198.

[43] L. A. Geddes and R. Roeder. "Criteria for the selection of materials for implanted electrodes". In: *Annals of Biomedical Engineering* 31 (7 2003), pp. 879–890.

[44] Yu Mike Chi, Tzyy Ping Jung, and Gert Cauwenberghs. "Dry-contact and noncontact biopotential electrodes: Methodological review". In: *IEEE Reviews in Biomedical Engineering* 3 (February 2010), pp. 106–119.

[45] Jong-hyeon Chang et al. "Fitting Improvement Using a New Electrical Circuit Model for the Electrode-Electrolyte Interface". In: (2007).

[46] L A Geddes. "Historical Evolution of Circuit Models for the Electrode-Electrolyte Interface". In: *Annals of Biomedical Engineering* 25 (1997), pp. 1–14.

[47] Wendy Franks et al. "Impedance characterization and modeling of electrodes for biomedical applications". In: *IEEE Transactions on Biomedical Engineering* 52 (7 July 2005), pp. 1295–1302.

[48] "Dry and noncontact EEG sensors for mobile brain-computer interfaces". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 20.2 (2012), pp. 228–235.

[49] Flurin Stauffer et al. "Skin Conformal Polymer Electrodes for Clinical ECG and EEG Recordings". In: *Advanced Healthcare Materials* 7.7 (2018), pp. 1–10.

[50] g.tec. *Scarabeo EEG Electrodes*. 2022. URL: `https://www.gtec.at/product/g-scarabeo-eeg-electrodes/`.

[51] Datwyler. *SSoftpulse EEG Electrodes*. 2022. URL: `https://datwyler.com/technology-innovation/advanced-technologies/strategic-fields/wearables-digital-health/softpulse`.

[52] g.tec. *Pangolin EEG Electrodes*. 2022. URL: `https://www.gtec.at/product/g-pangolin-electrodes/`.

[53] "Signal Comparison of Developed EEG Device and Emotiv Insight Based on Brainwave Characteristics Analysis". In: *Journal of Physics: Conference Series* 1505.1 (2020).

[54] "Towards conductive-gel-free electrodes: Understanding the wet electrode, semi-dry electrode and dry electrode-skin interface impedance using electrochemical impedance spectroscopy fitting". In: *Sensors and Actuators, B: Chemical* 277.August (2018), pp. 250–260.

[55] Guang Li Li et al. "Review of semi-dry electrodes for EEG recording". In: *Journal of Neural Engineering* 17.5 (2020).

[56] "A stretchable electrode array for non-invasive, skin-mounted measurement of electrocardiography (ECG), electromyography (EMG) and electroencephalography (EEG)". In: *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC'10* (2010), pp. 6405–6408.

[57] Seong Ho Yeon et al. "Acquisition of Surface EMG Using Flexible and Low-Profile Electrodes for Lower Extremity Neuroprosthetic Control". In: *IEEE Transactions on Medical Robotics and Bionics* 3.3 (2021), pp. 563–572.

[58] "Impedance sensing device enables early detection of pressure ulcers in vivo". In: *Nature Communications* 6 (2015), pp. 1–10.

[59] "A wearable biosensing system with in-sensor adaptive machine learning for hand gesture recognition". In: *Nature Electronics* 4.1 (2021), pp. 54–63. URL: http://dx.doi.org/10.1038/s41928-020-00510-8.

[60] Seung H. Ko et al. "All-inkjet-printed flexible electronics fabrication on a polymer substrate by low-temperature high-resolution selective laser sintering of metal nanoparticles". In: *Nanotechnology* 18.34 (2007).

[61] J. Chung et al. "In-tandem deposition and sintering of printed gold nanoparticle inks induced by continuous Gaussian laser irradiation". In: *Applied Physics A: Materials Science and Processing* 79.4-6 (2004), pp. 1259–1261.

[62] Jaewon Chung et al. "Conductor microstructures by laser curing of printed gold nanoparticle ink". In: *Applied Physics Letters* 84.5 (2004), pp. 801–803.

[63] P. Salvo et al. "A 3D printed dry electrode for ECG/EEG recording". In: *Sensors and Actuators, A: Physical* 174.1 (2012), pp. 96–102.

[64] "Micromachined electrodes for biopotential measurements". In: *Journal of Microelectromechanical Systems* 10.1 (2001), pp. 10–16.

[65] Ryan Kaveh et al. "Rapid and Scalable Fabrication of Low Impedance, 3D Dry Electrodes for Physiological Sensing". In: *Advanced Materials Technologies* (May 2022), p. 2200342.

[66] Mohammad Meraj Ghanbari et al. "A Sub-mm3 Ultrasonic Free-Floating Implant for Multi-Mote Neural Recording". In: *IEEE Journal of Solid-State Circuits* 54.11 (2019), pp. 3017–3030.

[67] Cheryl A Deckert. "Electroless copper plating. A review: Part I". In: *Plating and Surface Finishing* 82.3 (1995), pp. 58–64.

[68] Karthik Gopalan et al. "Vacuum Formed Coils for Magnetic Resonance Imaging". In: *2021 International Conference on Electromagnetics in Advanced Applications, ICEAA 2021* (2021), pp. 327–330.

[69] Wayne J. Staab et al. *A one-size disposable hearing aid is introduced*. 2000.

[70] Campbell. *Ear Illustrations*. 2019. URL: https://www.campbellmedicalillustration.com/.

[71] Aviral Pandey et al. "A 6.8 μW AFE for Ear EEG Recording with Simultaneous Impedance Measurement for Motion Artifact Cancellation". In: *Proceedings of the Custom Integrated Circuits Conference* 2022-April (2022).

[72] Benjamin C. Johnson et al. "An implantable 700uW 64-channel neuromodulation IC for simultaneous recording and stimulation with rapid artifact recovery". In: *Symposium on VLSI Circuits* (2017).

[73] Marco Guermandi, Eleonora Franchi Scarselli, and Roberto Guerrieri. "A Driving Right Leg Circuit (DgRL) for Improved Common Mode Rejection in Bio-Potential Acquisition Systems". In: *IEEE Transactions on Biomedical Circuits and Systems* 10 (2 Apr. 2016), pp. 507–517.

[74] B. Razavi. *Design of Analog CMOS Integrated Circuits*. 2nd ed. USA: McGraw-Hill Higher Education, 2016. ISBN: 9780077496128.

[75] Jiansheng Xu, Yisong Dai, and Derek Abbott. "A Complete Operational Amplifier Noise Model: Analysis and Measurement of Correlation Coefficient". In: *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.* 47.3 (2000), pp. 420–424.

[76] Jiawei Xu et al. "Measurement and analysis of current noise in chopper amplifiers". In: *IEEE J. Solid-State Circuits* 48.7 (2013), pp. 1575–1584.

[77] E. Huigen, A. Peper, and C. A. Grimbergen. "Investigation into the origin of the noise of surface electrodes". In: *Medical and Biological Engineering and Computing* 40 (3 2002), pp. 332–338.

[78] Arjang Hassibi et al. "Comprehensive study of noise processes in electrode electrolyte interfaces". In: *Journal of Applied Physics* 96 (2 2004), pp. 1074–1082.

[79] Walter G. Besio et al. "High-Frequency Oscillations Recorded on the Scalp of Patients With Epilepsy Using Tripolar Concentric Ring Electrodes". In: *IEEE Journal of Translational Engineering in Health and Medicine* 2.May 2013 (2014).

[80] Andy Zhou et al. "A wireless and artefact-free 128-channel neuromodulation device for closed-loop stimulation and recording in non-human primates". In: *Nature Biomedical Engineering* 3 (January 2019).

[81] Carolyn Schwendeman, Ryan Kaveh, and Rikky Muller. "Drowsiness Detection with Wireless, User-Generic, Dry Electrode Ear EEG". In: *2022 44th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2022).

[82] Thien Nguyen et al. "Utilization of a combined EEG/NIRS system to predict driver drowsiness". In: *Scientific Reports* 7.August 2016 (2017), pp. 1–10. ISSN: 20452322.

[83] Saskia Haegens et al. "Inter- and intra-individual variability in alpha peak frequency". In: *NeuroImage* 92 (2014), pp. 46–55.

[84] Terence W. Picton et al. "Human auditory steady-state responses: Respuestas auditivas de estado estable en humanos". In: *J. of Audiology* 42.4 (2003), pp. 177–219.

[85] Gary Rance et al. "Hearing Threshold Estimation in Infants Using Auditory Steady-State Responses". In: *J. of the American Academy of Audiology* 16.5 (2005), pp. 291–300. ISSN: 1050-0545.

[86] Jaehyuk Lee et al. "A 0.8-V 82.9-μ W In-Ear BCI Controller IC with 8.8 PEF EEG Instrumentation Amplifier and Wireless BAN Transceiver". In: *IEEE J. Solid-State Circuits* 54.4 (2019), pp. 1185–1195. ISSN: 00189200.

[87] Seunghyeok Hong et al. "Intelligent system for drowsiness recognition based on ear canal eeg with ppg and ecg". In: *Information Sciences* 453 (July 2018). DOI: 10.1016/j.ins.2018.04.003.

[88] F. Lotte et al. "A review of classification algorithms for EEG-based brain-computer interfaces". In: *Journal of Neural Engineering* 4 (2 June 2007).

[89] Deon Garrett et al. "Comparison of linear, nonlinear, and feature selection methods for EEG signal classification". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 11 (2 June 2003), pp. 141–144.

[90] Abdulhamit Subasi and Ergun Erçelebi. "Classification of EEG signals using neural network and logistic regression". In: *Computer Methods and Programs in Biomedicine* 78 (2 2005), pp. 87–99.