LOW COST NEUROCHAIRS

A Thesis

Presented to

the Faculty of California Polytechnic State University

San Luis Obispo

In Partial Fulfillment

of the Requirements for the Degree

Master of Science in Computer Science

by

Frankie Pike

December 2012

© 2012
Frankie Pike
ALL RIGHTS RESERVED

COMMITTEE MEMBERSHIP

TITLE: Low Cost NeuroChairs

AUTHOR: Frankie Pike

DATE SUBMITTED: December 2012

COMMITTEE CHAIR: Franz Kurfess, Ph.D.

COMMITTEE MEMBER: Alex Dekhtyar, Ph.D.

COMMITTEE MEMBER: Foaad Khosmood, Ph.D.

Abstract

Low Cost NeuroChairs

Frankie Pike

Electroencephalography (EEG) was formerly confined to clinical and research settings with the necessary hardware costing thousands of dollars. In the last five years a number of companies have produced simple electroencephalograms, priced below \$300 and available direct to consumers. These have stirred the imaginations of enthusiasts and brought the prospects of “thought-controlled” devices ever closer to reality. While these new devices were largely targeted at video games and toys, active research on enabling people suffering from debilitating diseases to control wheelchairs was being pursued. A number of neurochairs have come to fruition offering a truly hands-free mobility solution, but whether these results could be replicated with emerging low cost products, and thus become a viable option for more people is an open question. This thesis examines existing research in the field of EEG-based assistive technologies, puts current consumer-grade hardware to the test, and explores the possibility of a system designed from the ground up to be only a fraction of the cost of currently completed research prototypes.

Acknowledgements

This thesis was an incredible adventure and I would like to extend my appreciation to everyone who made it possible. CalPoly is an amazing learning environment and I can confidently say I learned a great deal from each professor from whom I had the pleasure of taking a class. The passion exhibited by the faculty here is incredible. My thesis advisor, Dr. Franz Kurfess, is owed special thanks for offering me such flexibility on this project, teaching me so much through a number of classes, and always leveraging his connections to enhance the quality of my education and maximize the success of this project. My friends added so much to the richness of my college experience as well.

For actually meeting with me, speaking so kindly, and loaning me hardware, I would also like to thank Stanley Yang, An Luo, and Jake Chuang from NeuroSky. NeuroSky is an incredibly interesting company with some amazingly talented engineers. I am very inspired by their work, passion, and products.

I had the opportunity to publicly demonstrate several core components of this project and want to thank everyone who took the time to stop by and provide feedback. In particular, the personal stories you shared about family and friends really affirmed my belief in these types of projects and made it easier to continue in the face of all manner of technical hurdles.

But perhaps most significantly, the on-going support of my family has been irreplaceable. I am forever indebted to them for their boundless encouragement, imparting belief in the value of education, and tolerating my take-it-apart nature from a young age. But not for their love and guidance, there is no way I would be where and even who I am today.

Contents

List of Tables	viii
List of Figures	ix
1 Introduction	1
2 Related Work	4
2.1 Alternative Wheelchair Control	4
2.2 Current Consumer-Grade EEG Hardware	5
2.3 EEG as an Input Modality	8
2.4 Noted Limitations of EEG	10
3 Design and Architecture	11
3.1 Prototyping Platform	11
3.2 Sensor Arrays	13
3.3 Multi-agent Architecture	19
3.4 System Decomposition and Integration	24
4 Experiments	26
4.1 The Emotiv Epoc	26
4.1.1 Participant Evaluation	27
4.1.2 Preparation for EEG Acquisition	28
4.1.3 Cognitiv Suite (EEG)	29
4.1.4 Verification of Trained Models	29
4.1.5 Findings	30
4.1.6 Expressiv Suite (EMG)	33
4.1.7 Form Factor Considerations	34

4.2	NeuroSky Devices	35
4.2.1	Alternative EEG Mechanisms	35
4.2.2	Hardware	38
4.2.3	The NeuroSky SDK	39
4.2.4	SSVEP Stimulator	39
4.2.5	Results	43
5	Contributions	48
6	Future Work	50
6.1	Novel Platforms	50
6.2	Enhanced DSP	51
6.3	Classifiers	51
6.4	Inter-Agent Cooperation	52
	Bibliography	53

List of Tables

4.1	Participant Trials	27
4.2	Participant Performance	30

List of Figures

2.1	Sample EEG Trace [3]	6
2.2	Emotiv Epoc [11]	8
2.3	NeuroSky MindWave [29]	8
3.1	Roomba 530 with Bluetooth Receiver	12
3.2	High Definition Camera Array	15
3.3	Edge Detection Prototype	16
3.4	Optical Flow Prototype	17
3.5	Stereo Correspondence Prototype	18
3.6	Stereo Correspondence Depth Map	18
3.7	Proposed Multi-agent Architecture	20
3.8	Bidding Agent Lifecycle	21
3.9	Auctioneer Lifecycle	21
3.10	Improved Auctioneer Agent Lifecycle	23
4.1	intendiX P300 Interface [5]	36
4.2	NeuroSky MindBand	38
4.3	SSVEP Stimulator (GDI+)	41
4.4	SSVEP Stimulator (SDL)	42
4.5	SSVEP Stimulator with EEG Processing	43
4.6	Harmonic Amplitudes at 11Hz	44
4.7	EEG Output Stability	44
4.8	Cross-Sensor Discrepancy at 11Hz	46

Chapter 1

Introduction

Diseases like amyotrophic lateral sclerosis (ALS, also known as Lou Gehrig's disease) affect the nerve cells in both the brain and spinal cord which are responsible for voluntary muscle control. While conditions like ALS carry with them a host of other health and psychological problems, there was one which struck me as surprisingly important: the loss of mobility associated with the progression of these diseases. Until seeing it taken from a family member, it had never fully registered how heartbreakingly significant mobility can be to an individual's self-efficacy. Electric wheelchairs can go a considerable way toward allaying some of these challenges, but when the fine motor skills required to control these devices wane, the user is left with little more than an expensive seat. Knowing there had to be a better alternative, I explored both commercially available products and novel control systems in research with the aim of at finding truly hands-free solutions.

There are a multitude of companies in the electric wheelchair market, with most catering to the elderly and selling in a direct-to-consumer fashion (though the end purchase is often subsidized heavily by health insurance providers). These "powerchairs" take two general form factors, one resembling a tricycle (controlled via a handlebar) and another more akin to a recliner with a joystick. The latter is better suited for ALS patients as it requires less physical exertion to guide.

Vendors like Walmart sell this form factor from around \$1,500 though features and branding can drive the cost well above the \$4,000 mark. Electric wheelchairs targeted at the medical industry can cost 17 times more, with a study from 2010, specifically focused on ALS patients, citing the average cost at \$26,404 [38]. To keep costs in perspective, the wheelchair alone is \$10,000 more expensive than the MSRP of a new Honda Civic. Companies like Adaptive Switch Laboratories offer modular control solutions for wheelchairs, extending their utility as conventional input modalities fall short. Pressure sensors along the headrest or a sip-and-puff decoder provide alternative control for users and can be replaced as needed. These systems appear to be representative of currently common industry offerings and add several thousand dollars to the total system cost.

This problem space struck me as incredibly interesting for a variety of reasons. First and foremost, advancements in this field can have a real and significant effect on the quality of life for patients. Bringing the force of technology to bear upon even minuscule parts of a person’s daily routine can transform struggle into independence. Second, commercially available solutions are paltry in comparison to the state-of-the-art in research. Far from a topic confined to small research groups, even big corporations like Toyota have demonstrated the viability of more exotic control systems like non-invasive electroencephalograph (EEG) headsets [36]. For patients under a variety of lock-in syndromes, these systems promise an incredible step forward in autonomy, unrealized by any currently available products. This thesis examines existing research in the field of EEG-based assistive technologies, puts current consumer-grade hardware to the test, and explores the possibility of a system designed from the ground up to be an order of magnitude more affordable than currently completed systems.

Such a wheelchair would have several core subsystems which are explored as independent components in this thesis. At the most basic level, there is the wheelchair itself with the necessary hardware bridge for decoding commands from a computer. A modified iRobot Roomba 530 was used as a prototyping platform since it is very compact while still embodying the drive mechanics of a conventional wheelchair. To bring the system to an actual wheelchair platform would

simply require translation at the point of the hardware-software interface. EEG processing is considered another subsystem (documented in chapter 4) since the neurological mechanisms exploited are tightly coupled to the capabilities and sensor configuration of the EEG device. Section 4.1 examines a procedure for using the Emotiv Epoc for brain-state detection and its associated difficulties while section 4.2 explores alternative possibilities with dry sensors on a NeuroSky MindBand. Both EEG solutions introduce uncertainty into the wheelchair’s reasoning since they serve as only probabilistic indicators of intent. To compensate for this property, a final subsystem is required which examines directives derived from EEG in light of situational data acquired from on board sensors. The multi-agent architecture responsible for environmental modeling and assuring the safe and continued operation of the wheelchair is the subject of chapter 3. Ideally, these components would have ultimately converged into a single cohesive prototype, but for reasons which are explained throughout the thesis, difficulties in harnessing the claimed capabilities of consumer products prevented the creation of a unified system.

Nonetheless, each of these subsystems represents a contribution toward the future realization a low cost NeuroChair. The state training methodologies developed with the Emotiv Epoc can be applied future projects, whether or not they depend on near-real-time control. The stimulator created for use with the NeuroSky MindBand is likewise transferable to future experiments involving the presence of steady state visual evoked potentials. The proposed multi-agent architecture is applicable to many other robotics projects and embodies a simple paradigm, yet is solidly grounded in past research. These contributions are chronicled in chapter 5 and chapter 6 highlights some of the many ways in which people from other disciplines could find worthwhile ways to expand on this thesis.

Chapter 2

Related Work

2.1 Alternative Wheelchair Control

One of the earlier applications of EEG to wheelchairs is known as The “Hands-free” Wheelchair Control System (HaWCoS) [13]. While prior work seemed to focus on applications of electro-oculography (EOG, allows estimation of eye position) or the detection of slow cortical potentials (SCPs), the HaWCoS used electromyography (EMG) to detect volitional contraction of facial muscles. These electrical impulses were then interpreted by an on-board computer and used to command the wheelchair. Though not exactly an example of the brain computer interfacing (BCI) realized today, it did show the viability of a hands-free system and established a meaningful metric for comparison with traditional control mechanisms. As part of their experimental validation, the notion of time “overhead” was used. Predetermined paths have a known optimal time based on the maximum speed of the wheelchair, which participant trial runs can be compared against. Impressively, participants with The HaWCoS required only about a 50% overhead, affirming the utility of their platform for practical use.

Since the HaWCoS, other EEG mechanisms apart from SCPs have been tested for their viability to control computer systems. Some of these are discussed in

section 4.2.1, but the most compelling implementation was called the Bremen Autonomous Wheelchair Rolland [26]. This platform coupled an EEG cap detecting steady-state visual evoked potentials (SSVEP) with a laser range finder to evaluate the safety of any command. In experimental settings, SSVEP-based BCI has been shown to provide an acceptable level of usability on tasks related to selection and navigation [18] and even been successfully combined with VEP/P300 (another type of visually stimulated BCI modality) wheelchair systems [32]. The Rolland also showed that SSVEP is largely unaffected by participant stress levels (i.e. being on a moving wheelchair) and provided classification accuracies greater than 90% [26]. Since the Rolland was equipped with a laser ranger, localization was core to its navigation, and users selected among computed paths. For structured environments, like homes, this can greatly simplify the demands placed on the user. For example, they may only need to select among a list of rooms and have the unit handle the navigation on its own. The HaWCoS also explored this possibility but noted that, particularly in cases of assistive wheelchairs, there is also an implicit need to ensure that users feel as though they are *driving* the wheelchair as opposed to simply *riding* along with a computer [13]. Laser rangers impart substantial situational awareness to the platform, but also several thousand dollars to the total system cost. This thesis attempted to use the same neurological mechanism as the Rolland, but couple it with a significantly cheaper array of high definition web cameras, infrared sensors and ultrasonic rangers to see if similar success is achievable at a lower cost.

2.2 Current Consumer-Grade EEG Hardware

Electroencephalograms, whether consumer or medical grade, are intended to measure electrical activity produced by the interaction of neurons within the brain. While these devices are quite sensitive (often using conductive gel to minimize distortion caused by electrical wiring in the surrounding room) the electrical impulses of a neuron are far too small to be detected individually. Yet the brain contains billions of these neurons, and their interactions give rise to “ionic gra-

Figure 2.1: Sample EEG Trace [3]

dients” across neuronal membranes [8]. These gradients can propagate like a wave toward the scalp and create a *push* or *pull* on the electroencephalogram’s conductive sensors at the scale of 1 to 20 microvolts. These signals are then amplified, and when logged over time, produce EEG traces as seen in Figure 2.1. Software development kits (SDKs) for consumer products tend to compute a power spectrum for these waveforms and collapse these into meta-signals based on their frequencies for analysis. Waves between 8 and 12 Hertz are known as alpha waves and are associated with being mentally alert, but with eyes closed. Beta waves have frequencies greater than 13 Hertz and are associated with being alert or actively focusing. Delta waves, with frequencies less than 3 Hertz, and theta waves, between 4 and 7 Hertz, and are both associated with sleep. Taken together these provide certain insight into the subject’s current brain state and possibly mood. Most devices also support some method of access to raw sensor data which can provide resolution down to the level of a single Hertz (rather than clustered meta-signals). These capabilities were explored in this thesis to extract other interesting neurological phenomena in section 4.2.1.

While EEG clearly arises from the medical and research fields, the dawn of consumer applications for this technology has, perhaps curiously, arisen in the computer gaming sector. Quality-of-life enhancements based on this technology may be underway, but gamers have already had commercially available EEG-powered products since 2007 [1]. There are several big players in the field, notably

NeuroSky, Emotiv, and OCZ Technology, who have developed an impressive array of products. CalPoly is fortunate enough to have several devices from these companies available for student projects, including the Emotiv Epoc and the NeuroSky MindWave.

The Emotiv Epoc (Figure 2.2) has an array of 16 saline sensors and was heavily marketed toward the video game industry. It is a dual purpose device, capable of both electroencephalography (EEG, electrical activity associated with neurons) and electromyography (EMG, electrical activity associated with skeletal muscles). This facilitates the detection of not only brain states, but also facial expressions and eye blinks. Various components of the software development kit are also targeted at detecting certain emotional states like excitement and frustration. The consumer version of the headset currently costs \$299 while a research kit (which includes the needed software development kit) costs \$750, making for a low cost unit with an impressive set of capabilities. The positions of sensors on the Epoc correspond to a subset of those used in the International 10-20 system, though in clinical settings, 21 electrodes are recommended [2].

NeuroSky has a wider selection of products, ranging from the \$99 MindWave (Figure 2.3) to prebuilt components for inclusion in other products. NeuroSky chipsets provided the EEG capabilities of Mattel’s Mindflex and Uncle Milton’s Star Wars Force Trainer which both enjoyed impressive commercial success. Their products generally consist of a single dry sensor placed on the forehead, providing the benefit of simplicity at the expense of functionality [28]. The MindWave and MindSet are of this form factor and monitor a user’s degree of focus, though unlike the Emotiv Epoc, EMG capabilities and the training of brain state models are absent. NeuroSky also offers proprietary development kits, centered on specific use cases (like the detection of steady state visual evoked potentials, SSVEP) on a case-by-case basis. I was fortunate enough to make contact with Dr. Ann Luo, NeuroSky’s manager of algorithm development, and borrow additional hardware for this thesis. At a corporate level, NeuroSky’s focus seems to be on the development of components, which are ultimately brought to market, or given compelling use by partners [27].

Various products have been released by other companies, but were not explored as part of this thesis. Market successes have largely been confined to toys, as evidenced by NeuroSky’s latest creation, a pair of EEG responsive cat ears, dubbed Nekomimi. Though each of these products has cultivated a loyal following of enthusiasts, the proverbial “killer app” for consumer-grade devices appears to have yet to emerge.

Figure 2.2: Emotiv Epoc [11]

Figure 2.3: NeuroSky MindWave [29]

2.3 EEG as an Input Modality

A recent hardware survey examined currently available BCI control mechanisms for wheelchair use and, like this thesis, placed an emphasis on low-cost solutions [33]. Both the Emotiv Epoc and NeuroSky MindSet made an appearance in its analysis as indicators of how quickly embedded signal processing has advanced. The author also expressed a bit of surprise that the ever falling cost of EEG acquisition devices, coupled with their targeting to the gaming sector, has not made the seemingly small step from controlling virtual entities on a screen to guiding a tangible object for the benefit of users. A number of neurological mechanisms were examined for their applicability to wheelchair control, reliability, and ease of implementation. SSVEP, the same mechanism used by Bremen’s Rolland [26], hybrid systems [32], and the subject of a new NeuroSky development kit [25], was chosen as the best candidate, lending further justification for

its use in this thesis. The survey stopped short of implementation rendering this thesis an extension and validation of its conclusions.

As with any unconventional input modality, it is also important to consider usability in addition to its mere feasibility. Even though the proposed end-user is likely to be limited in their alternatives, there is no reason to diminish efforts to make the system operate as naturally as possible. EEG is far from mainstream adoption, but there have fortunately been a number of recent studies aimed at refining the definition of usability in the context of brain computer interfacing [18] [6] [24]. As a developer, it is easy to focus on information transfer rate (ITR), classifier reliability, and state discrimination times since these provide concrete metrics on the performance of the system. Nonetheless, users inherently come in with their own sets of expectations for the system, and perceptions of system responsiveness can be just as important to the value of the system. When applied in formats like brain state mapping, which attempt to identify trained “thoughts”, the user is put in the situation of seemingly needing to convince the device of what they are thinking. Systems which purport to “read” minds fall short and can generate frustration. For virtual interactions or experiments this is undesirable, but in order to seriously consider EEG as an input modality for frequently used devices like wheelchairs, these faults must be addressed. Furthermore, unlike keyboards, mice, or joysticks, which provide users with instantaneous tactile feedback, BCI has a detached component of processing to determine user intent. Such dilemmas are not unique to BCI and have correlates in other input methods like speech recognition. Interestingly, when users were given the opportunity to perform the same tasks with a speech recognizer and an EEG apparatus there were only minor differences in the reported usability [18]. The BCI version was noted as more difficult to learn, but scored higher in terms of user enjoyment, and lower on the metric of “tiring”. The researchers attribute this primarily to BCI’s novelty, but even so, these are certainly worthy qualities to bring to a system aimed at increasing a user’s independence.

2.4 Noted Limitations of EEG

While there have been some remarkable successes in the creation of research prototypes, there has also been some harsh criticism targeted at the use of EEG as a primary input source, especially those dependent on slow cortical potentials (SCP). The prospects for the technology are quite exciting, but early experimenters were dismissive of its capability as a real-time solution [13][4]. Even in clinical settings with expensive collection equipment, artifacts induced by muscle movement (i.e. blinking) or electrical interference can have a significant impact on the quality of the signal captured.

Beyond just generic problems with EEG, there have also been considerable concerns expressed over practical implementations as applied to wheelchair control. Unlike headrest sensors, which have definitive on and off states, current EEG methods do not exhibit this characteristic. To correct for errors in detection, most methods conduct multiple consecutive trials in an aim to reduce false positives. While increasing the accuracy of the system, this also limits the information transmission rate (ITR) of current implementations [35]. Hybrids of multiple collection methods have been explored to play on comparative strengths and weaknesses [32], but even these fail to provide full fail-safety for the operation of a wheelchair. A potential solution is the coupling an EEG control system with a partially autonomous mode aimed at assuring the continued, safe operation of the system [20][26]. The interplay between these two systems becomes an interesting engineering effort, unlike most instances where user commands can be considered authoritative. Fortunately, the field of autonomous robotics is well explored, and multi-agent systems and partially observable Markov decision processes (POMDPs) show potential as solid architectures for this class of problem [30][39].

Chapter 3

Design and Architecture

3.1 Prototyping Platform

As previously noted, the use of EEG as an exclusive input method can be prone to error caused by a wide variety of conditions, ranging from the distraction of the operator to variations in the local electrical field. Lower cost EEG devices can only compound some of these concerns, introducing critical issues for assuring the safe operation of the wheelchair. While able-bodied individuals would be able to press a conventional override button, the target population for this system may not have this option, and baking such functionality into the EEG system would fail to resolve this problem independent of its likely source. Remote overrides can be provided to caretakers, but in order to facilitate both independent control by a paralyzed user and autonomous navigation modes an on-board sensor array is needed.

It is easy to conceive of naïve, monolithic control systems where disparate inputs are amalgamated for a single decision engine. The engine is responsible for weighing various heuristics and choosing the correct course-of-action based on the entire environmental state. Such systems have achieved considerable success, and one need not look much further than the iRobot corporation's Roomba

Figure 3.1: Roomba 530 with Bluetooth Receiver

product line. Arrays of simple sensors and scripted behaviors can clean rooms surprisingly well and appear to exhibit at least some form of intelligence within the operating domain. In implementation, the facade of intelligence fades away to stratified behaviors, where the system chooses among ever more sophisticated behaviors as lower-level conditions (like safe distance from drops) are assured. The use of the Roomba Serial Communication Interface (SCI) allows the disabling of tiered safety checks, rendering the unit almost humorously unintelligent. While crashing into walls and charging off stairs can serve to make prototyping more light-hearted, it highlights a crucial shortcoming of such dependency on tiered behavior: breakdown at lower levels can render higher order planning meaningless. These issues cannot be completely mitigated, as ultimately any robotic system inherently relies on sensor data, but in cases like wheelchair control, expected behaviors do not lend themselves easily to the same type of stratification which has proven successful in the Roomba. Incidentally, the Roomba features similar drive mechanics to electric wheelchairs, and ended up serving as the prototyping platform for this thesis.

3.2 Sensor Arrays

Being so heavily explored, one of the initial challenges was identifying system architectures that worked well, rather than those whose further work sections alluded there was much to be desired. The cost constraint also represents another challenge since many wheelchair systems feature \$5,000+ laser ranging systems. Alternative, low cost sensors have the consequence of increasing the degree of uncertainty in the platform’s internal world representation, also requiring models which are flexible in reasoning amid uncertainty. Partially observable Markov decision processes, an extension of Markov chains, are aptly suited to this task, though their flexibility comes at a price: they are PSPACE-hard [31]. Nonetheless, approximations can be used to reduce the scope of the state distributions and bring the benefits into computational tractability.

Laser ranging systems coupled with localization algorithms have proven themselves to be quite capable on the wheelchair platform, particularly in well defined areas [26][14]. A considerable portion of the wheelchair’s time is likely to be spent in a home, hospital, or similar environments where map-based solutions can be quite effective. Even if the goal is not a fully or partially autonomous system, a safety sensor array can still serve as a helpful override for errant EEG readings. On the other hand, due to the greater demands placed on a user for asserting control via EEG, some users may find it very desirable to select among common destinations in their home, freeing themselves to interact with guests rather than exclusively focusing on the task of navigation. Regardless of which system a user would prefer, laser rangars can add a significant amount to the total cost. Many projects coming out of the campus robotics club used incredibly cheap components for collision avoidance (infrared and ultrasonic sensors) and commodity webcams with OpenCV for path planning.

Though it was not a wheelchair, the RHINO also calls to attention important considerations for sensor arrays on autonomous systems [14]. While relying on a conventional array of sensors (stereo cameras, SONAR, and laser ranging), RHINO was situated in a museum for testing, giving it certain obstacles which

were “invisible”, at least to its sensors. Glass casings around exhibits were essentially undetectable and created discrepancies between SONAR and laser readings which might otherwise seem like an errant situation. The use of Markov localization allowed RHINO to nonetheless navigate the environment. Taking into account largely “invisible” obstructions which a wheelchair may encounter (i.e. chain-link fencing or glass panes next to doors) will be key to the user experience and underscores the need to ensure no single sensor, no matter how expensive it is, is afforded too much control over the system. In the case of glass walls, a 150 times cheaper ultrasonic sensor ended up being correct, rather than the higher resolution laser.

While laser rangefinders (in tandem with other sensors) are still likely the best solution, part of this project aimed to identify whether or not a “good enough” solution might be possible for under \$100 (less than 10% of the cost of a commonly used SICK laser). Android smartphones, for example, frequently have an impressive array of sensors, which could bring the platform not only a persistent Internet connection for mapping related functions, but also a great degree of situational awareness. Digital compasses would be clearly beneficial for orientation and GPS could be utilized outdoors for routing. Better yet, if the user happens to own one of these devices, integration with the computer on which the BCI software is running is simple and affords added navigation capabilities at no cost. Unfortunately, experimentation with several higher end devices (Samsung Galaxy S2, Google Nexus, and a Qualcomm development phone) showed their compass capabilities to be woefully unreliable on moving platforms. Immediately after calibration, they were placed on a Roomba which was programmed to query the phone for its current heading and navigate a 2 meter by 2 meter square. Distances were handled by internal odometry, but the points at which turns finished were dictated by compass values. None of the phones were suited to this task, and while internal odometry could have been used to handle turn timing, after retrieving the phones and comparing their readings when pointed to their original start position, all had at least 20 degrees of deviation. Though the compasses may prove problematic, outdoors, GPS appeared to be consistently reliable on

campus and heading was able to be determined over the course of several seconds of GPS readings.

Figure 3.2: High Definition Camera Array

For longer range obstacle detection, this thesis experimented with a variety of methods using high definition web cameras (at a total cost of just under \$60) and an open source collection of image processing libraries called OpenCV. The first prototype simply tested the ability of a single camera to identify edges without any particular concern to ranging. This is a fundamental task to allow the platform to travel through corridors with minimal user intervention. Joysticks provide users with greater turning precision, but when using EEG the number of commands which can be deciphered is more limited. The natural language command “proceed straight down the hallway”, which users would reasonably expect to be executable, carries an implicit command that straight is relative to the corridor and not the device’s current bearing. A user who approached a corridor at an off angle, equipped with only left and right commands, would have a significantly less enjoyable experience navigating such a device versus one capable of small corrections. Though less significant when navigating across a campus, most any indoor environment will feature corridors, making this a significant usability enhancement. Figure 3.3 shows the prototype identifying edges of a potential obstacle.

Figure 3.3: Edge Detection Prototype

Transitioning from simple edge detection to corridor navigation is a matter of applying the same techniques and identifying changes over time. A common solution is based on “optical flow” wherein the same points are tracked across multiple frames, giving these points motion vectors, and the platform a means of visual odometry. Assuming interesting points are found on both sides of the hallway, a program can determine whether it is proceeding approximately down the center or converging on either of the two sides by examining the magnitudes of the motion vectors. If the platform is converging on a wall, the magnitudes will be unbalanced, with smaller values on the side being converged upon. Figure 3.4 shows the optical flow prototype traveling down the center of a hall. The points in green are those being tracked across frames and the center line shows the division of points into the left and right fields. The carat in the top left text indicates the device intends to continue traveling straight since the flow estimates 0.82 and 0.88 (for the left and right fields respectively) are within the tolerance threshold. On the 1.5 GHz Atom netbook used for testing, this task is able to run approximately 5 times per second allowing the device to make quick corrections.

Although optical flow techniques work well for corridor navigation, they in-

Figure 3.4: Optical Flow Prototype

herently suffer from the inability to identify obstacles or recommend bearing adjustments when the platform is stopped. The biological answer to this problem is frequently depth perception via stereo vision, which is achievable with OpenCV as well. Using a pair of high definition webcams, affixed to a piece of PVC, the third prototype calibrates to the disparity between the two cameras and correlates the resulting pixel differences with depth. Similar indoor robotics projects seem to prefer a Microsoft Kinect which provides this functionality (along with many others) through their SDK. Unfortunately, the Kinect's abilities are greatly diminished outside since it relies on the detection of an IR projection rather than stereo correspondence. Wheelchairs are apt to be taken outside, making this limitation rather significant. Figure 3.5 illustrates a calibrated output from the two cameras with Figure 3.6 as the resulting depth map. Depth maps range from black to white (in grayscale) with white representing the closest objects. It is not quite the output of the Kinect, but the detection of part of the wall gives hope that further refinements could render a \$60 camera configuration into a very capable obstacle detector. Fortunately, stereo correspondence and optical flow are not mutually exclusive. Both can operate concurrently, or be selectively enabled

and disabled (perhaps based on platform speed) to conserve compute resources.

Figure 3.5: Stereo Correspondence Prototype

Figure 3.6: Stereo Correspondence Depth Map

While the Roomba has its own sensors for detecting drops and wall proximity, dedicated ultrasonic sensors (connected to an Arduino) were also briefly explored. When mounted to the front of a modified remote control car, the sensor was able to successfully alert the on-board micro-controller and affect a stop. Due to their lack of precision (which increases as a function of distance) these would predominately be used as a safety override to prevent imminent collision, rather than

localization as lasers might otherwise be. As the most significant component of this thesis is its EEG integration these prototypes were not taken to completion. Nonetheless, they do show the plausibility of low-cost sensor arrays as an alternative to the higher resolution (and more importantly higher cost) systems more commonly seen in research platforms.

3.3 Multi-agent Architecture

In the case of the wheelchairs, one of the great challenges in implementing a monolithic control system is difficulties in acquiring sensor results, and consequently formulating a representation of the system’s state. In particular, those related to computer vision can return at varying and unpredictable intervals unlike simpler queries to ultrasonic sensors on a micro-controller. Return times from vision processing can take long enough as to invalidate other readings, and GPS presents the challenge of varying accuracy, where the most recent reading may not, in fact, be the best to consider. These limitations are of direct relevance to the system’s internal representation of the environment and can lead to situations of stalled action and uncertainty. For instance, if a GPS fix is lost or has been reduced to an unusable accuracy (20+ meters), it is uncertain whether it is best to wait for a new position to be acquired (which may be a fraction of a second away), or rely on a previous reading from some arbitrary time in the past, perhaps corrected with internal odometry. Whichever option is chosen, the prevailing concern is one of bringing the system to a halt based on uncertainty at any given tier of decision making. To design a more robust system, this thesis opted for an architecture based on agents with delegated responsibilities and an implementation of partially observable Markov decision processes to help mitigate uncertainty in reasoning.

To aid in the reduction of the system’s belief space (the number of possible states it may be in), the proposed architecture utilizes agent consensus. Rather than defining thousands of potential states and their associated transition functions (as may be done in a more formal Markov decision process), these can be

abstracted into higher order goals on a per-agent basis. The relative weighting of goals encodes the reward function of the Markov decision processes and becomes one of the core goals in development. Reasoning through various contentions (like short term planning agents with their longer term counterparts) is unlikely to follow a single heuristic and, more likely, will be context dependent. This context dependency embodies the finite set of actions and available transitions from each state. In this manner, consensus driven multi-agent systems can then be used to make the modeling of Markov decision processes substantially easier for a developer, allowing emphasis to be placed on higher level interactions than the minutia of state transitions. The use of multiple agents can also afford a degree of isolation and compartmentalization to the system further easing development.

Figure 3.7: Proposed Multi-agent Architecture

Figure 3.7 illustrates the categories of agents in the proposed system and some of their delegated responsibilities. NavClasses in this case is a general placeholder for classes which will contain the logic necessary to receive percepts from the sensor arrays and extract meaningful data from them. NavClasses are called by independent agents as needed to arrive at a decision. Likely targets are shown below the individual bidding agents, like the Obstacle Avoidance System’s (OAS) use of 3D cameras to extract depth information and estimate time-to-collisions. The auctioneer operates as a consumer of bid notifications until notified by the auction agent that the current auction has ended. The listed agents are by no means exhaustive and are limited only by the imagination of developers and availability of sensors.

Implementation of the individual agents is quite simple in this framework. The tasks that would otherwise make this program more complex, like synchronization, are already handled by the agent library. This leaves only the core logic and agent lifecycles up to the developer.

Figure 3.8: Bidding Agent Lifecycle

Figure 3.9: Auctioneer Lifecycle

Bidding agents have a simple task at a high level. They call necessary methods in the NavClasses, evaluate the results, and submit a bid. Interestingly though, in addition to competing with the other bidding agents, agents can also potentially compete with themselves. This is a consequence of the high refresh rate of certain navigation components. GPS positions are relatively easy to process, and can be updated several times each second, affording the agent multiple opportunities to bid against itself and handle fluctuating fix quality. Depending on the auction's duration, subsystems with high latencies (like 3D video processing) may not even

bid each round. This too can be desirable since it allows the system as a whole to continue functioning despite what might have otherwise been a bottleneck.

The auctioneer also follows a cyclic pattern wherein it starts auctions, identifies high bids as they are received, and waits for a time notification from the auction agent. This does not stop agent bidding. In fact, frequently in simulations, bids are still received after the timer ends, though they are disregarded from the current auction's result. This may strain the auction metaphor, since the incoming bids can be higher than the auction's best and received only a millisecond late. At some point though it is necessary to terminate the auction and take action, lest the robotic platform wait indefinitely for a "better bid" which may never come. Since the auctions are constantly cycling, in reality, these "missed bids" will simply be regarded in the next auction, preserving the value of the agents' efforts. After the auction is ended, the highest bid's directive is sent to the actuator controller and the next auction is started.

In simulations, it became clear that early versions which treated auctions as independent decision making events suffered from what might externally be viewed as indecisiveness. Auction results, in a worst-case scenario, can yield functionally negating command sequences (i.e. LEFT, RIGHT, LEFT, RIGHT...). This is clearly unacceptable in a practical sense, and suggests the need for some degree of persistence on the level of the auctioneer to safeguard against such results. A decision should not be favored indefinitely, but should exert some force on the directives which follow. This is implemented as memory with an associated decay function.

The modified process, illustrated in Figure 3.10, stores the last directive that resulted from an auction and its associated bid. This bid, in effect, becomes the next minimum bid with which agents must compete. This greatly reduces the prevalence of L/R/L/R sequences by requiring a higher degree of confidence on the part of the bidding agent. Following the earlier model, the agent merely needed a hunch. The decay function serves as a countermeasure to what could be considered the opposite problem of indecisiveness: stubbornness. Without decaying the stored value, it is possible for a high confidence assertion to exert

Figure 3.10: Improved Auctioneer Agent Lifecycle

protracted control over the system, even when the very subsystem which produced it no longer holds the claim. Confidence values are in the range 0-100, and the decay rate is set at -20% per auction. With auctions taking place approximately once per second, this limits a percept with 100% confidence to exerting control over the system for 5 seconds, worst case scenario. This seems acceptable, but fortunately, this decay function can easily be fine tuned to become proportionate to the confidence value, change depending on the assertion, or incorporate other emergent concerns. Nonetheless, the storing of decisions and its associated decay methodology does seem to provide a more consistent output.

As a matter of implementation, the first version of this architecture was created using the Microsoft Asynchronous Agent Library (AAL) which comes bundled with Visual Studio 2010. Rather than having to handle thread management, locking, unlocking, and deadlocks, this agent platform allows the developer to focus more on logic than concurrency. The nature of a timed auction also brings the system closer to “real time” though it does not satisfy the more rigid definition. Having the various subsystems, proxied by agents, asynchronously reporting is ideal and keeps the decision engine running on a continually fresh set of percepts. Unfortunately, the Microsoft Asynchronous Agent Library introduces a platform dependency and limits the ability to bring the end system to potentially novel platforms like cell phones. The second version of the agent architecture used the user-defined message passing capabilities of Simple DirectMedia Layer (SDL)

giving it a greater degree of platform independence. Though it does not affect the platform as implemented, message passing between agents is more difficult since the intended paradigm in SDL is the consumption of all messages by the primary thread of execution. Since the auctioneer is spawned in that thread, the current system holds, but efforts to enable agent cooperation may be more difficult under SDL than the Microsoft AAL.

3.4 System Decomposition and Integration

To aid in its development, a NeuroChair can be decomposed into several connected subsystems: the physical platform and sensors, the EEG components, and the decision engine which mediates between them. This thesis explored each subsystem independently with the aim of ultimately integrating them into a single system upon independent validation.

The hardware layer, representing what would ultimately be the wheelchair and safety sensor arrays was modeled through the use of an iRobot Roomba 530. Powered wheelchairs feature the same types of drive mechanics, making the Roomba an effective prototyping platform with significantly greater portability. The use of the Roomba serial communication interface (SCI) over Bluetooth allowed full control over the platform and retrieval of the the Roomba’s own on-board bump and drop sensors. The PVC-mounted high definition cameras and ultrasonic sensors were the intended deployment solutions, regardless of whether they were mounted to a Roomba or wheelchair. As a platform, the Roomba was abstracted within the software. Possible commands were *stop*, *forward*, *turn 90 degrees left*, *turn 90 degrees right*, or *forward* with a 45 degree deviation left or right. These were translated into platform specific commands by an intermediary layer, rendering the migration to other chairs possible without significant revisions to the architecture as a whole. Such abstraction is likely to be helpful in supporting multiple wheelchair models in the future as well. Cameras and other sensors were similarly replaceable.

The subsystems responsible for EEG processing are outlined in the next chapter. Much like the abstraction layer for the base platform, applying the same paradigm to EEG hardware confers future flexibility in terms of device form factor and capability sets. Different types of sensors and neurological mechanisms can be explored with only minor modifications to the mediating architecture since they share a common, simple interface.

The use of EEG rather than conventional input devices like joysticks introduces uncertainty into the control system. User intent, rather than having binary discernibility (a button is either depressed or inactive), will only be expressed through correlations with other states. To reason amid this uncertainty, the multi-agent architecture proposed in this chapter was implemented. This subsystem provides the critical bridge between the platform and the EEG hardware, allowing commands derived for EEG and cameras to be validated against safety sensors. The multi-agent architecture is designed to be fault-tolerant albeit with an impacted safety margin. Depending on the sensor set ultimately chosen for a full-scale prototype, there would need to be established thresholds for when the system should no longer be able to proceed if certain sensors no longer return data.

Though each subsystem must be validated independently, it is also important to analyze the performance and integrity of the system upon integration. The decay values proposed in the auction agent’s life cycle must be empirically determined based on the output patterns of the device, lest these amount to little more than “magic numbers”. Concrete usability assessments can be achieved by comparing paths traveled between a participants equipped with a NeuroChair and others on a standard power chair, as done with the HaWCoS [13]. Finally, a balance must be identified between the reward functions for user commands and bids by agents using sensor data. Unfortunately, such experiments depend on EEG components to be part of an integrated system which was not achieved in this thesis. Consequently, validation of the intended system remains pending a successful EEG command acquisition system.

Chapter 4

Experiments

4.1 The Emotiv Epoc

The first headset explored in this thesis was the Emotiv Epoc. Its software development kit provides a wealth of APIs to target EEG, EMG, and even gyroscopic data from the device. While encrypted for lesser editions of the SDK, the Research Edition also provides access to raw EEG, allowing for the development of even more sophisticated algorithms. EEG-centric activities are collected into the *Cognitiv* suite, under which I built several prototypes.

Promotional material for the Epoc, and even a TED talk given by Emotiv Lifescience’s CEO Tan Le [23], routinely use the word *thought* when explaining what the unit detects. In reality, that is an exceptionally high level of abstraction, which may give users some false expectations. The Cognitiv suite performs machine learning on user-tagged data, allowing it to run a classifier on future inputs. I, along with my test subjects, expected to be able to ultimately map thoughts like *left*, *right*, and *stop* onto control inputs for the device. Indeed, during the TED talk, an individual is brought out, performs a single session of training, and within a few seconds is able to reproduce their mapped state to the satisfaction of the classifier, all within a span of 2 minutes. This would be in line with expec-

Participant	Time (s)	Appx. Hours
1	20,523	6
2	6,920	2
3	3,178	1
4	11,655	3
5	9,254	2

Table 4.1: Seconds of EEG logged

tations of a purportedly *thought-controlled* wheelchair, but unfortunately, such an implementation proved to be elusive.

4.1.1 Participant Evaluation

In an attempt to find a reproducible methodology for training the device a number of participants were brought in for several hours each. A pre-built application (Control Center) which comes bundled with the SDK allowed for simple experimentation and provided the first glimpse into potential difficulties. It allows users to train a cognitive model in 8 second intervals, assess how closely new training data matches the existing model, and interactively test their ability to reproduce a given brain state. Both intuition and Emotiv’s documentation suggest that the reliability of the models increases over time. Consequently, for this evaluation it made most sense not to attempt a large number of short experiments, but to conduct extended sessions. The training process is quite involved and requires a significant time investment before the data that will ultimately be required is even logged. Table 4.1 breaks down the amount of EEG data logged within the SDK for each participant.

Surprisingly, participants were already aware of the Emotiv Epoc, even if they are unacquainted with the fact that it is an EEG device, and were quite eager to take part in experimentation. This made the introduction relatively straightforward, though there were a few common misconceptions to dispel, largely for the sake of participant comfort (namely that it cannot *read* minds or in any way tamper with the brain). Participants were also assisted when putting on

the device to ensure that the sensors align with their reference positions. The Control Center is then consulted to ensure that each electrode has an acceptable connection to the participant’s scalp.

4.1.2 Preparation for EEG Acquisition

Based on preliminary experimentation with the Epoc and posts on their development forums, it seemed that mapping arbitrary *thoughts* was an exercise in futility. It remained to be seen whether or not my concurrent roles as experimenter and subject were the primary source of difficulty, so the first participants were given very little direction, allowing them to map states as they saw fit. After approximately 30 minutes with little success, they were then instructed to avoid both abstract and concrete thoughts and focus on more kinesthetic actions. This instruction proved helpful (evidenced by classifier matches becoming non-zero) and this advice was given to the remaining participants at the beginning of testing. While it is generally advisable to avoid priming research participants, even the Emotiv development forums suggest this tip to avoid frustration and speed up training sessions.

As a matter of procedure, before Cognitiv state training can begin there is some initial work required to prepare a user profile. The participant’s unique neutral state must be trained before proceeding further. The Control Center has a mode which allows for protracted recording in 30 second intervals. To attempt to keep this data as realistic as possible, the subject is not informed of the significance of these recordings, but is conversed with as normal. Initial experimentation showed that focusing too actively on being calm during neutral recording resulted in erratic behavior later. This may be a result of an implied imperative to avoid thoughts pertaining to actions that will later be trained. Keeping the participant distracted should also build a more realistic model, since the end product is not intended to be operated in a quiet, isolated environment, but become a part of the user’s daily routine.

4.1.3 Cognitiv Suite (EEG)

Emotiv’s Control Center gives states physical metaphors like *push* and *pull*. The goal was to have subjects iteratively train 4 states mapping to *left*, *right*, *stop*, and *continue*. While this may seem like coarse control for a robotic platform, discrimination between 4 states pushes the capabilities of the Epoc and is an upper limit imposed by the SDK.

To acquaint them with the process, participants performed 3 rounds of 8 second training on the *left* action. They were then shown a screen with a centered box and asked to reproduce the left action. If it is successfully detected, the box moves in response. This gives participants a great deal of excitement, and through discussion we attempt to identify the critical portion of their *thought* to the detection. For example, with one subject it was the point at which the energy they envisioned as traveling down their arm arrived at their hand. This gives the subject a better idea of how the system responds and hopefully speeds up the actual training. Once the participant is acquainted with the process of training and validation, their data logged up to this point (with the exception of their neutral data) is reset. This gives the training an effective fresh start, free of possible contamination by earlier exercises. The actions *left*, *right*, *stop*, and *continue* were then trained iteratively.

4.1.4 Verification of Trained Models

Though the described training procedure already requires participants to exhibit a degree of control over a Cognitiv state, these are validated in a serial fashion which is incongruous with the intended deployment environment. While it is critical to ensure that the recorded state is reproducible at the time of logging, the real test of a user’s mastery must be the issuance of commands on demand. To prevent users from claiming credit for effectively random triggers, participants were to be presented with a sequence of randomly chosen directives on index cards and timed for how long it took them to signal the state. The

Participant	Left	Right	Stop	Continue
1	12%	4%	0%	0%
2	34%	14%	0%	0%
3	8%	2%	0%	0%
4	65%	28%	4%	0%
5	27%	10%	0%	0%

Table 4.2: Match rates for participants

triggered sequence of Cognitiv states was also to be recorded to measure deviations. Unfortunately, participants were unable to master the training process, rendering this testing protocol moot. The few times participants were prompted in this manner, excessive periods of time (greater than 30 seconds) were spent waiting for any state to be triggered. For the intended purpose such a delay is simply unacceptable.

As part of its control center, the SDK also provides estimates for the reproducibility of Cognitiv states. These percentages do operate in a bit of a black box manner, and the formula for their direct calculation is not provided in documentation. Nonetheless, the documentation does at least indicate they are measures of the correlation between repeated training sessions of the same action. They begin at 0 when the subject has only trained once and, ideally, increase with the number of training sessions and a greater convergence of the training data. Observationally, it was also noted that the greatest increases in these scores occurred when the action receiving supplemental trained was being detected simultaneously (based on previous recordings). Table 4.2 outlines the performance of participants.

4.1.5 Findings

It is difficult to identify the root cause of the difficulties encountered. Based on the discrepancies between participants, it is likely that the success of state training depends greatly upon the person. Nonetheless, it also became clear that the process used for training is of consequence as well. There are many

threads on the Emotiv support forums dedicated to Cognitiv training protocols, though being user submitted (rather than company verified) they should be taken with a grain of salt. Many of these personal experiences seemed to be good starting points, but experimentation did reveal some general principles that were consistent across test subjects.

Poor things to train on

Words: This is unfortunate, partly because if it did work this way, the interface would be incredibly natural. Promises of “mind control” carry with it the assumption that the device will simply accept thoughts which are easily made manifest as words. In reality, the Epoc is more adept at identifying feeling. No amount of training while simply envisioning or internally vocalizing words seemed to work for any participant.

Abstract Concepts: Though these seem relatively close to words, it is commonly held in marketing that colors can also carry a charged emotional context. There was some hope that envisioning colors or similarly abstract concepts like “friendship” might carry enough associated emotion to build a reliable model. In practice, this was not the case.

Concrete Objects: To get away from the other categories, subjects also frequently fixated on concrete objects, like a stop sign (intuitively mapped to stopping the device), or even more specific instances like the stop sign on their walk to class. This is probably still too high a level of abstraction to be reasonably detected via EEG signals (or at least those acquired by devices like the Epoc).

Better things to train on

Emotions: Perhaps comically, this was discovered after several hours of trying to get the device to respond to anything. After briefly admitting defeat, the device started triggering. Apparently frustration was the most salient

portion of the signature, not the intended action. While initially surprising, this does make sense considering the Epoc SDK has an entire section dedicated to emotional state estimation. It also advises some caution for ensuring participants are in a relatively calm state during training, lest it interfere with the model being built.

Kinesthetic actions: It now comes as little surprise that the SDK’s training suite is so insistent on using kinetic metaphors for actions, rather than simply user defined fields. Participants noted that it was substantially easier to hold onto a single thought longer (which is significant since training sessions are 8 seconds long) when it involved some mental “exertion”. Envisioning themselves pushing or pulling something worked well and is likely a good place to start with future participants.

Meditative Activities: Akin to yoga or meditation, some participants attempted breathing activities and focusing on individual sections of their body with a non-zero degree of success. It also likely helps to avoid the emotional contamination that proved difficult initially.

There was not a universal trigger that worked well across participants, but these guidelines collectively did aid in the training process and allow them to gain a degree of control over the system. The results of this experiment were disheartening, but certainly not conclusive. There are, however, still some practical concerns which are worth noting. Even dismissing the time necessary for initial setup, the time required to train a sophisticated model is nowhere near the 5 minute long Emotiv demo. Any successful implementation of this method will undoubtedly require a patient user and an understanding that the Cognitiv models take a great deal of refinement. Admittedly, this experiment did aim to push the capabilities of the device by using the maximum number of trainable states within the SDK. In Emotiv’s defense, their live demo was predominantly an instance of single state detection, for which one participant did receive a rating of 65%. Unfortunately the degree of success in reproducing more than 2 Cognitiv states was essentially zero. Repeated training was of little assistance and the

more effort that was focused on training later states, the more difficult it became to reproduce the first two states. Fatigue and frustration may well have been a contributor to the declining abilities of the subjects, but these will undoubtedly be the sentiments of frustrated end users as well.

Commentary on collaborative EEG projects does note that the Emotiv product has notably poor signal impedance relative to “medical/research-grade” EEG devices [12] and Emotiv, themselves, insist the device is not intended for medical purposes. While this thesis is unconcerned with diagnoses, the devices used in previous work were frequently medical-grade, yielding perhaps unrealistic expectations for what is, in fact, consumer hardware. In spite of the claimed limitations, there are still plenty of videos on the Internet showing people exhibiting an impressive degree of control via the device and it may just be a matter which requires an even greater time investment. As a practical matter, however, the greatest concern for this thesis is the ability to issue directives reliably as there are clear safety implications. Even extending the greatest observed control across all 4 actions, active control 65% of the time is still unacceptable for a primary input device.

4.1.6 Expressiv Suite (EMG)

In addition to its processing of EEG, the Emotiv SDK also has logic for distinguishing among various facial gestures, interpreted from EMG signals. Blinks, brow furrowing, eye direction, and various other facial operations could also be used as a control mechanism. While perhaps less intuitive than what EEG promises (“just think it”), these actions have the benefit of being far more controllable, and with practice, could become as second-nature as using a joystick to control a wheelchair. False positives for events like eye blinks could be controlled for by deliberate contrivance (i.e. requiring a user to blink their eyes out of phase), in a way that EEG monitoring may be unable to match. The gyroscopic sensor in the Epoc could also be creatively leveraged for verification.

A prototype was developed with this mechanism, though this functionality is

akin to the HaWCoS [13], rather than highlighting the EEG capabilities of the Epoc. Generic EMG signatures ship with the Emotiv SDK making for an easier demo (since it does not require the state training of EEG), but this version was not actively pursued since it has already been shown to be viable in past work and requires a greater degree of volitional muscle control.

4.1.7 Form Factor Considerations

When using the Epoc for extended periods of time (> 2 hours), the reported signal quality gradually degrades. To restore signal quality, it is recommended that the user take off the headset, remove the contact assemblies, and re-moisten the felt pads with saline solution. The process takes only a few minutes, but does bring up a potential usability concern for the Epoc. Able-bodied participants could easily be trained to complete this process (which does not preclude the Epoc from a wide variety of uses), but the target population of this thesis may not be in such a position. Technologies aimed at promoting independence must also be concerned with similar human factors, something which appears to be beyond the intended purpose of the Epoc. For now, the need to rehydrate the contacts represents a serious usability concern.

Emotiv forum posts dating back several years suggest that dry contacts were in active development for the Epoc which could go a significant way toward mitigating this concern. To date, these contacts have yet to be made available to consumers, though posts in 2010 alluded to their near-term availability [10]. An inquiry to their technical staff yielded claims that development has gone well, but their time-to-market is unknown. The proposed contacts are designed to be a drop-in replacement for existing units, which could still make the Epoc a worthwhile headset to reevaluate in the future and increase its utility for assistive technologies.

4.2 NeuroSky Devices

After considering the results from the Emotiv Epoc and the more significant human factors issues, it seemed appropriate to explore EEG devices equipped with “dry” sensors. These sensors can be worn for hours at a time, thereby increasing the independence with which a potential user could operate the device. NeuroSky has several devices with dry sensors, but the capabilities of their products lag behind the feature set of the Emotiv Epoc. First and foremost, while the Emotiv has 16 sensors (still less than the recommended 21 for lab usage [2]), NeuroSky devices generally only feature 1 or 2 [28]. This limits their ability to isolate signals from different regions of the brain (for example, targeting areas associated with visual processing) or collect EMG artifacts apart from eye blinks. Nonetheless, the simplicity of the form factor has a great deal of practical appeal and could reduce the overall complexity of the wheelchair system.

4.2.1 Alternative EEG Mechanisms

While an EEG signal is accessible via the NeuroSky hardware, developing a state-mapping solution seemed to be a less compelling direction if the better situated sensors of the Epoc were having issues performing reliably. While investigating alternatives another class of solution showed great potential. Rather than attempting to map entire states to actions, other methods focus on predictable, cross-subject-consistent physiological reactions to visual stimuli.

VEP/P300 (Visually-Evoked Potentials), a type of event related potential (ERP), was the first candidate considered. With a sensor positioned near the brain’s parietal or occipital regions, a user is asked to pay attention to an item on a screen. Targets on the screen, typically arranged in columns and rows are highlighted in an arbitrary sequence, eventually hitting the target the user is focused on. While there are slight variations across subjects, a small deflection in voltage (known as an “oddball response”) is detectable approximately 250 to 600 milliseconds after the onset of an infrequent stimulus. Systems with multi-

ple targets perform a deductive process to determine which target the user was focused on by performing multiple trials. If there are many targets, switching between highlighting columns and rows is also used. The Emotiv Epoc has sensors at both O1 and O2 and has been applied successfully to the collection of P300 responses [9]. Though the researcher noted difficulty in the collection process, P300 responses have also been collected using the NeuroSky MindBand with a participant success rate of 78.5% [16]. P300 detection is included in the popular OpenViBE package and has been performed using a wide variety of custom EEG solutions [22]. There are commercial applications of this technology, like the gtec intendiX SPELLER (Figure 4.1), which allow users to interact with virtual keyboards, though it requires a more sophisticated setup with a dedicated signal processor, caps for reference positions, and gel for the electrodes [17]. Nonetheless, VEP/P300 appears to be a well-trodden path with considerable success, even on consumer-grade hardware. The lack of intense per-user training has great appeal, especially in light of gtec’s claim of 10 minutes between introduction and successful use [17].

Figure 4.1: intendiX P300 Interface [5]

Though not a safety issue for keyboard emulating systems, the greatest concern with VEP/P300 in the context of this thesis is its relatively low information transfer rate (ITR). The intendiX, for instance, claims users are generally able to enter 5 to 10 characters per minute. While certainly impressive from a technical perspective, this suggests an average command time between 6 and 12 seconds,

which on a moving platform could lead to situations which compromise user safety (at a modest 3 MPH a wheelchair would cover just over 16 meters). For comparison, simulated experiments using SSVEP for wheelchair control were able to achieve command discrimination within 2 seconds (allowing the wheelchair to travel only 2.6 meters at the same velocity) [35]. For this project fewer targets (as opposed to a full QWERTY keyboard) could be used, increasing the ITR. Research in P300 performance has acknowledged the inherent trade-off between the accuracy imparted by multiple trials and the need for higher ITRs and sought to reduce the cost, but gains appear to be relatively modest compared to the ITRs achieved by other BCI methods [34].

SSVEP (Steady State Visually Evoked Potentials) detection is similar in spirit to VEP/P300, using visual stimuli to trigger a detectable reaction in the brain, but is also a good candidate for applications requiring a higher ITR. While ERPs (like P300) can be thought of as a reaction to a certain visual stimulus, SSVEP are an indicator of the presence of a certain visual stimulus. Targets oscillating at frequencies below 75 Hz can trigger synchronous (and resonant) responses near the observer’s occipital lobe. While the Epoc has sensors already in position, devices like the NeuroSky MindBand can also be worn in reverse to position their 1 or 2 sensors in this area for detection. Unlike VEP/P300, the events which need to be detected are not time-locked making their detection slightly easier. The need to have targets oscillating at distinct frequencies limits the number of commands which the user can select between at any given time, but in the context of a wheelchair control system, 4 to 6 candidates seems reasonable, especially if targets can switch purpose based on context. To determine if a user is attending to a given target (i.e. one oscillating at 11 Hz) the system examines the user’s EEG power spectrum at 11 Hz, 22 Hz, 33 Hz, etc. against a defined threshold. From a processing standpoint, this is less demanding than the support vector machines which facilitate the discrimination of brain states, bringing computational requirements down to the level of smartphones [37]. Like P300-based BCI, very little per-user customization is required and when coupled with the promise of a higher ITR, SSVEP seems like an ideal candidate for

exploration.

4.2.2 Hardware

While a research team had performed ERP detection with the MindBand [16], the use of single or dual sensor MindBands for SSVEP seemed to be untested. A paper published by two researchers at NeuroSky centered on SSVEP detection algorithms and alluded to the creation of a new, capacitive (non-contact) sensor [25]. These sensors have a great appeal due to their ability to mitigate the effects of hair which otherwise degrades signal quality. Probes to the NeuroSky sales team indicated that the product was not ready for purchase, though their website references an SSVEP development kit for select partners. I was able to get in contact with Stanley Yang (NeuroSky’s CEO) and Dr. An Luo (one of the authors of the paper and NeuroSky’s Manager of Algorithm Development) who were very supportive of this project. A dual sensor MindBand (Figure 4.2) was loaned to me for experimentation in this thesis, though the device does not appear to be the unit used in the NeuroSky paper.

Figure 4.2: NeuroSky MindBand

4.2.3 The NeuroSky SDK

The NeuroSky SDK comes in a variety of versions which allow developers to target a number of different platforms. In addition to the major operating systems, there are also packages for smartphones and even small scale devices like the Arduino Duemilanove. As a result of having significantly fewer sensors than products like the Epoc, the capabilities exposed by the SDK are more limited, allowing the tracking of a few meta-signals like attention along with the raw EEG values. The lack of pre-built functionality for the detection of SSVEP allowed this thesis to simply use the SDK's stream parser, with some minor modifications to accommodate the hardware's unique protocol. The data received from the device is unencrypted (unlike the Epoc), allowing applications to communicate with the device over Bluetooth using a standard serial port profile.

Unlike the Emotiv SDK which provides access to a power spectrum from the device, NeuroSky's SDK provided only power band values. The distinction between theta (4-7 Hz), beta (13-30 Hz), and gamma (31-50 Hz) bands may be sufficient for many applications, but capturing SSVEP requires buckets at the level of a single hertz. This was accomplished by obtaining the power spectrum via the raw signal passed through KissFFT. NeuroSky providing their SDK as C source made this exchange very straightforward.

4.2.4 SSVEP Stimulator

In order to detect SSVEP, a tightly controlled visual stimulus must be created. OpenViBE has a SSVEP scenario, but collecting the required dependencies to build the stimulator proved difficult. The NeuroSky MindBand used in this thesis also had a slightly different communication protocol which is currently unsupported by OpenViBE, so there was incentive to build a more portable, stand-alone application which could also be free of constraints imposed by a larger, more generic BCI package. A previous project developed an open source SSVEP stimulator based on OpenGL and provided a great starting point for this

portion of the project [7].

Stand-alone LED stimulators, like the one used in the NeuroSky paper [25], allow flexible timing and targeting of frequencies known to elicit stronger responses (the 9-13 Hz range), but computer based solutions have additional limits imposed by the inherent refresh rates of monitors [40]. Fine-grained timing can be further impacted by additional loads on the system. Four frequencies were chosen for detection (6, 10, 12, and 15 Hz) which fall evenly within a 60 Hz refresh rate and have only 1 harmonic collision (the first harmonic of 6, and the stimulus 12). Though an external stimulator would have brought increased flexibility, the proposed architecture for the control system already includes a laptop, and using its display as the stimulator (in addition to providing insight into the application status) reduces the overall complexity of the implementation.

The low cost netbook used in prototyping did not feature a dedicated graphics coprocessor and the integrated graphics chipset supported an outdated implementation of OpenGL. This caused issues for the open source stimulator when it attempted to initialize. On modern hardware there were no issues, but to decrease the barrier to entry, I developed two alternative versions which did not have this restriction.

The first version (Figure 4.3) was a rudimentary implementation with Microsoft's Graphics Device Interface (GDI+), a predecessor to Direct2D. This version executed without complaint on the netbook, but was unable to take advantage of the more accurate timing OpenGL afforded the other implementation. Additionally, when the system came under load by other processes (even on newer hardware) there was a noticeable fluctuation in the output refresh rate. Unfortunately, this could not be mitigated through the use of a graphics coprocessor on newer hardware (since GDI+ relies exclusively on software rendering) and even minor changes in output yield frequencies which are no longer the target of detection.

DirectX's newer interfaces support acceleration, but the most enticing feature of the OpenGL version is its cross-platform compatibility. Various libraries with this characteristic were examined, and Simple DirectMedia Layer (SDL) provided

Figure 4.3: SSVEP Stimulator (GDI+)

an excellent set of capabilities across a wide variety of platforms [21]. When the system is capable, SDL can take advantage of hardware acceleration and preload textures to graphics memory. Though the animations for this application are simple alternating checkerboards, the rate at which they flicker can benefit from high resolution timers and optimized blitting operations. The second prototype (Figure 4.4) utilized SDL and performed consistently well on both the netbook and higher performance laptop.

With the stimulator achieving stable FPS counts, the next step was integrating the EEG processing components into the program. Simple graphing functions were also created to allow real-time insight into the EEG components. Figure 4.5 shows the completed SSVEP stimulator while it is actively acquiring data from the headset. Between the familiar checkerboards are two graphs. The top set of slanted blue lines is used to indicate disruptions in signal acquisition. Data packets from the device have a timestamp ranging from 0 to 255 (resetting on overflow) allowing the program to determine when data is lost. Since this can

Figure 4.4: SSVEP Stimulator (SDL)

have a significant effect on the computed power spectrum, the system actively monitors packet loss and only performs a Fast Fourier Transform (FFT) in the presence of a completed dataset.

The second graph shows the power spectrum from the most recent FFT. Detection frequencies and their harmonics are color coded for quick viewing. The red peak on the right side of the graph corresponds to 60Hz, where the mains (power line) hum should be detected. Based on the refresh rate of the device, a DFT which provides the 1Hz buckets necessary for discrimination between the 6, 10, 12, and 15Hz target frequencies could only be executed once a second. To give the unit a finer resolution, the program executes FFTs from a sliding window where some data points are shared between consecutive FFTs. This appears to be a normal practice and is adjustable within the application, with a default of 2 epochs per second.

Figure 4.5: SSVEP Stimulator with EEG Processing

4.2.5 Results

After the FFT implementation appeared to be completed, the first thing needing to be verified was the degree to which harmonics resembled each other. Existing systems tend to combine the fundamental detection frequency with its first 2 or 3 harmonics. In theory, they should bear some resemblance, especially if a user is actively focusing on the target. Figure 4.6 represents approximately 15 seconds of EEG activity.

It was very encouraging to observe the general similarities between the fundamental frequency (11Hz) and its subsequent harmonics. It is also worthwhile to note that the deflections become increasingly pronounced in higher harmonics (in particular around the 23rd epoch). Unfortunately inter-harmonic stability is only part of the story.

The main difficulty encountered in SSVEP detection was attempting to find

Figure 4.6: Harmonic Amplitudes at 11Hz

significant events in outputs which fluctuated wildly. To aid in later analysis, the program creates running CSV files broken down by frequency. Figure 4.7 shows a particularly egregious 20 second sample from the application which would make simple detection methods based on thresholds completely invalid.

Figure 4.7: EEG Output Stability

The target frequencies (coupled with their harmonics) appear to be at least loosely covariant, but this sample is clearly a departure from the relatively stable readings graphed in Figure 4.6. An interesting phenomenon, also included in the graph, is the fluctuation at 60Hz. This particular frequency (at least in the United States) should be dominated by the mains hum and be a bit more stable

than the lower frequencies, but curiously, they are not. These fluctuations may be expected on a moving platform, as it travels through various rooms, but both collections were taken with the participant in a stationary position. This dilemma led to several different methods of attempting to stabilize the incoming data.

The first method was an attempt to stabilize relative to the mains hum. Assuming 60Hz should be relatively stable and the detection frequencies are covariant, the program can attempt to identify a correction factor based on fluctuations at 60Hz to apply to the amplitudes at other frequencies. While notch filters are usually implemented to correct for distortion due to mains, this occurs at the SDK level which is bypassed in this implementation. Unfortunately, this naïve, yet simple solution only made amplitude fluctuations less pronounced and does not address potential root causes noted in other papers (like electrical fluctuations in the local environment, unrelated brain processes, and noise from the electrodes themselves [26]).

There were also two methods aimed at exploiting the MindBand’s multiple sensors. The Rolland wheelchair creates a linear combination of the values from multiple sensors in an effort to reduce the *nuisance signal* [26]. Figure 4.7 actually contains such a combination of sensors, but whether or not they were combined the same types of fluctuations were observed. Treating one sensor’s output as a noise channel and subtracting it from the other also was unhelpful, for reasons which are apparent in Figure 4.8.

For a neurological phenomenon which has an concentrated area of occurrence, there are surprisingly few times when the two sensors agree on the amplitude at any given frequency. Linear combinations of these sensors amidst wide disagreement seems like it should have very little impact on this issue. Using one as a difference channel is similarly haphazard. Handling this cross-sensor difference (and the likely related problem output instability) was the major impediment to further progress since this dilemma did not appear to be addressed in the documentation of similar projects. As suggested in [15], there are several possible methods for combining signals from multiple sensors, but whether the proposed methods resolve situations like this was left unaddressed.

Figure 4.8: Cross-Sensor Discrepancy at 11Hz

In a final attempt, the open-source libSVM package was integrated into a prototype in an attempt to potentially resolve this issue through the use of a multi-class support vector machine (SVM). Rather than using linear classification based on thresholds, SVMs are able to project the points into higher-dimensional space and identify hyperplane dividers. Fortunately, libSVM bundles a fair amount of this functionality making the core responsibility one of classification. This required the development of a training component for the stimulator. Similar to the process used with the Epoc, users can be asked to fixate on a given target for some number of seconds while the application logs their EEG. The computed amplitudes for each target frequency, along with the amplitude at 60Hz were provided to the SVM as a 5-dimensional data point, labeled according to the frequency being trained. Fortunately for participants, building this data set is easier than with the Emotiv SDK since it relies on the detection of an involuntary neurological response rather than requiring sustained cognitive effort. After each target has been trained, the SVM then computed discrimination planes allowing later inputs to be classified. libSVM's tremendous flexibility is a mixed blessing since it necessitates a large number of user-specified parameters and selectable kernel functions. Successfully applying machine learning to this task would likely be a thesis in itself and attempts to select among recommended, general purpose settings did not provide an appreciable increase in control. Though an SVM adds another layer of complexity, once implemented, it is likely to be more robust than

the arbitrary thresholds defined in other implementations.

Despite sincere and varied efforts, target discrimination with the MindBand proved to be elusive. Nothing suggested this problem is insurmountable, and greater expertise in the field of digital signal processing and machine learning, along with a more intimate knowledge of the device’s technical limitations might one day make short work of this current block. Nonetheless, for this thesis, the viability of the NeuroSky MindBand for wheelchair guidance is still unknown. A paper published by researchers at NeuroSky suggests an alternate SSVEP detection methodology, based on the time domain rather than the frequency domain [25]. While this paper used more sophisticated (and more expensive) hardware, the process may still be transferable to products like the MindBand. However, the greatest benefits of their SLIC method center on the detection of irregular stimuli which is not a concern in this use case. Being time-locked, their method requires precise timing from the external stimulator which is less reliable with a software stimulator. The stimuli flicker rates were also specifically chosen based on past research and their coincidence with a monitor refresh, limiting the ability to explore other target frequencies.

Chapter 5

Contributions

While a full-scale prototype may be yet unrealized, this project did have its fair share of contributions. Though in testing subjects were unable to master the full capabilities of the Emotiv Epoc, there was still some useful insight gained for training protocols in later projects. The SSVEP stimulator developed for the NeuroSky experiments could be quite useful for future projects or even provided as an alternative for the current solution in OpenViBE (which failed to work on either of my computers). Like [7], it would be great to make this thesis' code base open source since it utilizes an alternative graphics library, already has integration with kissFFT and libSVM, and operates within a multi-agent architecture.

New sensors are supposedly on the horizon from both Emotiv and NeuroSky which could drastically modify the outcomes of this thesis [10][25]. These new sensors will better mitigate the effects of hair, electrical interference and have better signal to noise ratios. Even when these make it to market, this thesis can still serve as a helpful guidepost, pointing out the types of difficulties to expect and possible solutions.

Finally, the implemented multi-agent architecture has great applicability to projects in robotics, whether they feature EEG or not. By design, it has significant strengths in the areas of fault tolerance, performance, and extensibility. Developers looking for a solid cross-platform architecture for autonomous robots

should strongly consider a similar implementation.

Chapter 6

Future Work

This thesis was thoroughly interesting to work on, and has clear challenges for people coming from a variety of disciplines moving forward. While I wish I could have explored each of these in depth myself, doing so would have gone well beyond the scope of a single thesis. Nonetheless, there is still a great deal to explore. Those looking to aid in the creation of a low cost system should find the following areas to provide rewarding technical challenges.

6.1 Novel Platforms

An entailed goal of this project was to keep processing requirements within capabilities of lower cost hardware. A dual core netbook was purchased for benchmarking, since its form factor affords several key benefits: low-cost, small footprint, and an extended battery life. If this product were to be deployed to an actual wheelchair, space and power requirements may become a more significant concern. Stereoscopic correction via OpenCV became prohibitively slow, averaging about 7 frames per second (FPS), compared to 50+ FPS on a higher performance laptop. While in some applications this drop may be acceptable, longer delays cause difficulty in cross-frame motion estimation, which was used in one

the prototypes. Later development of the SSVEP stimulator also highlighted the difficulty low-cost netbooks have with high display frame rates.

Apart from waiting for the steady march of faster low budget computers, the next generation of portable gaming devices may prove to be an interesting target for a feasibility study. Devices like the Sony PlayStation Vita incorporate multi-core processors, a dedicated graphics coprocessor, a GPS sensor, a digital compass, tilt and roll sensors, and pervasive Internet connections. A Qualcomm development phone borrowed for the thesis had similar capabilities which are already being applied to BCI [37]. Individuals with an interest in computing architectures could easily begin the transition from “normal” computers to devices offering significant consolidation, portability, and cost-savings.

6.2 Enhanced DSP

The digital signals processing used in this thesis was relatively rudimentary. Recent research, [19] for instance, centers on the minimization of spectral leakage to increase the quality of the power density estimation which is critical to the successful detection of SSVEP. Individuals with expertise or interest in spectral analysis, filters, or DSP in general have a great deal to add to this process. The quality of the power spectrum has a direct bearing upon the entire system and would be a great area for further refinement.

6.3 Classifiers

A relatively simple support vector machine (SVM) was implemented, though its output left much to be desired. Emotiv’s CEO noted that the software powering the Epoc relies on machine learning, but did not dive into the specifics of their implementation [23]. Given the successes of their product, even an SSVEP-based solution may benefit from further exploration. As noted in 4.2 the many types of kernel functions and their associated parameters, applied to BCI, could be an

interesting subject of both research and experimentation. Individuals interested in machine learning and data analysis could make some significant advancement on this front, increasing the robustness of the system as a whole.

6.4 Inter-Agent Cooperation

The version of the multi-agent architecture implemented in this thesis exposes just a tiny portion of the possibilities. One of the most compelling additions may be the creation of a generic mechanism for inter-agent cooperation. Isolated agents (responsible for their own sensors) work for simple platforms, but it is easy to imagine sensor meshes, across a variety of devices which could provide supplemental information to the wheelchair’s guidance system. The isolation of agents confers a degree of fault-tolerance, but agent cooperation could preserve this resilience, while still allowing agents to share percepts. Near field communication (NFC)-enabled tiles, QR tags, and WiFi networks could orient the wheelchair in foreign environments and even help multiple chairs interact with each other. Use cases are limited only by ingenuity and bound to capture the interest of those in the field of robotics and artificial intelligence.

Bibliography

- [1] Mind games. *The Economist*, March 2007.
- [2] American Electroencephalographic Society. Guideline 7: A proposal for standard montages to be used in clinical EEG, 1994.
- [3] Biomedresearches. Epilepsy Awareness Program - EEG Resources. http://www.biomedresearches.com/root/pages/researches/epilepsy/eeg_resources.html, March 2012.
- [4] N. Birbaumer, N. Ghanayim, T. Hinterberger, I. Iversen, B. Kotchoubey, A. Kuebler, J. Perelmouter, E. Taub, and H. Flor. A spelling device for the paralyzed. *Nature*, pages 297–298, 1999.
- [5] D. Boros. intendiX User Interface. March 2011.
- [6] D. P.-O. Bos, M. Poel, and A. Nijholt. A study in user-centered design and evaluation of mental tasks for BCI. In *Proceedings of the 17th international conference on Advances in multimedia modeling - Volume Part II*, MMM’11, pages 122–134, Berlin, Heidelberg, 2011. Springer-Verlag.
- [7] J. Boyd and Y. Chen. An open source stimulator for SSVEP-based BCIs. In *Proceedings of the 50th Annual Southeast Regional Conference*, ACM-SE ’12, pages 124–129, New York, NY, USA, 2012. ACM.
- [8] M. Dr. Selim Benbadis, M. Dr. Aatif Husain, M. Dr. Peter Kaplan, and D. Dr. William Tatum. *Handbook of EEG Interpretation*. Springer Demos Medic Series. Demos Medical Publishing, 2007.

- [9] H. Ekanayake. P300 and Emotiv Epoc: Does Emotiv Epoc capture real EEG? August 2010.
- [10] Emotiv Systems Inc. Epoc forums: Dry sensor availability. <http://www.emotiv.com/forum/forum12/topic302/>, February 2010.
- [11] Emotiv Systems Inc. Emotiv media. <http://www.emotiv.com/corporate/media.php>, November 2012.
- [12] S. Fairclough. Emotiv Epoc and the triple dilemma of early adoption. <http://www.physiologicalcomputing.net/?p=1191>, December 2010.
- [13] T. Felzer and B. Freisleben. HaWCoS: the "hands-free" wheelchair control system. In *Proceedings of the fifth international ACM conference on Assistive technologies*, Assets '02, pages 127–134, New York, NY, USA, 2002. ACM.
- [14] D. Fox. *Markov Localization: A Probabilistic Framework for Mobile Robot Localization and Navigation*. PhD thesis, University of Bonn, Germany, 1998.
- [15] O. Friman, I. Volosyak, and A. Graser. Multiple channel detection of steady-state visual evoked potentials for brain-computer interfaces. *Biomedical Engineering, IEEE Transactions on*, 54(4):742–750, April 2007.
- [16] M. Grierson and C. Kiefer. Better brain interfacing for the masses: progress in event-related potential detection using commercial brain computer interfaces. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '11, pages 1681–1686, New York, NY, USA, 2011. ACM.
- [17] g.tec Medical Engineering. intendiX: worlds first personal BCI speller. http://www.biosignal.at/Intendix/intendix_speller.htm.
- [18] H. Gürkök, G. Hakvoort, M. Poel, and A. Nijholt. User expectations and experiences of a speech and thought controlled computer game. In *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology*, ACE '11, pages 53:1–53:6, New York, NY, USA, 2011. ACM.

- [19] G. Huang, J. Meng, D. Zhang, and X. Zhu. Window function for EEG power density estimation and its application in SSVEP-based BCIs. In *Proceedings of the 4th international conference on Intelligent Robotics and Applications - Volume Part II*, ICIRA'11, pages 135–144, Berlin, Heidelberg, 2011. Springer-Verlag.
- [20] I. Iturrate, S. Member, J. M. Antelis, A. Kbler, and J. Minguez. A non-invasive brain-actuated wheelchair based on a P300 neurophysiological protocol and automated navigation. *IEEE Trans. on Robotics*, pages 614–627, 2009.
- [21] S. Lantinga. Simple DirectMedia Layer. <http://www.libsdl.org/>, June 2012.
- [22] lbonnet. OpenViBE P300 Speller. <http://openvibe.inria.fr/openvibe-p300-speller/>, September 2011.
- [23] T. Le. A headset that reads your brainwaves, July 2010.
- [24] F. Lotte. Brain-computer interfaces for 3D games: hype or hope? In *Proceedings of the 6th International Conference on Foundations of Digital Games*, FDG '11, pages 325–327, New York, NY, USA, 2011. ACM.
- [25] A. Luo and T. J. Sullivan. A user-friendly SSVEP-based brain-computer interface using a time-domain classifier. *Journal of Neural Engineering*, 7(2):026010, 2010.
- [26] C. Mandel, T. Lüth, T. Laue, T. Röfer, A. Gräser, and B. Krieg-Brückner. Navigating a smart wheelchair with a brain-computer interface interpreting steady-state visual evoked potentials. In *Proceedings of the 2009 IEEE/RSJ international conference on Intelligent robots and systems*, IROS'09, pages 1118–1125, Piscataway, NJ, USA, 2009. IEEE Press.
- [27] NeuroSky Inc. NeuroSky chipsets. <http://neurosky.com/Business/ThinkGearChipsets.aspx>.
- [28] NeuroSky Inc. NeuroSky technologies: Grounded in medical research. <http://company.neurosky.com/bci-technology-application-march>.

- [29] NeuroSky Inc. NeuroSky MindWave. <http://store.neurosky.com/products/mindwave-1>, November 2012.
- [30] Y. Ono, H. Uchiyama, and W. Potter. A mobile robot for corridor navigation: a multi-agent approach. In *Proceedings of the 42nd annual Southeast regional conference*, ACM-SE 42, pages 379–384, New York, NY, USA, 2004. ACM.
- [31] C. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Math. Oper. Res.*, 12(3):441–450, Aug. 1987.
- [32] B. Rebsamen, E. Burdet, Q. Zeng, H. Zhang, M. Ang, C. L. Teo, C. Guan, and C. Laugier. Hybrid P300 and mu-beta brain computer interface to operate a brain controlled wheelchair. In *Proceedings of the 2nd International Convention on Rehabilitation Engineering & Assistive Technology*, iCRE-ATe ’08, pages 51–55, Kaki Bukit TechPark II., Singapore, 2008. Singapore Therapeutic, Assistive & Rehabilitative Technologies (START) Centre.
- [33] K. Stamps and Y. Hamam. Towards inexpensive BCI control for wheelchair navigation in the enabled environment - a hardware survey. In *Proceedings of the 2010 international conference on Brain informatics*, BI’10, pages 336–345, Berlin, Heidelberg, 2010. Springer-Verlag.
- [34] H. Takahashi, T. Yoshikawa, and T. Furuhashi. Reliability-based automatic repeat request with error potential-based error correction for improving P300 speller performance. In *Proceedings of the 17th international conference on Neural information processing: models and applications - Volume Part II*, ICONIP’10, pages 50–57, Berlin, Heidelberg, 2010. Springer-Verlag.
- [35] A. Teymourian, T. Lueth, A. Graeser, T. Felzer, and R. Nordmann. Brain-controlled finite state machine for wheelchair navigation. In *Proceedings of the 10th international ACM SIGACCESS conference on Computers and accessibility*, Assets ’08, pages 257–258, New York, NY, USA, 2008. ACM.
- [36] Toyota Central R&D Labs Inc. and Genesis Research Institute Inc. Real-time control of wheelchairs with brain waves. June 2009.

- [37] Y.-T. Wang, Y. Wang, and T.-P. Jung. A cell-phone based brain-computer interface for communication in daily life. In *Proceedings of the 2010 international conference on Artificial intelligence and computational intelligence: Part II*, AICI'10, pages 233–240, Berlin, Heidelberg, 2010. Springer-Verlag.
- [38] A. L. Ward, M. Sanjak, K. Duffy, E. Bravver, N. Williams, M. Nichols, and B. R. Brooks. Power wheelchair prescription, utilization, satisfaction, and cost for patients with Amyotrophic Lateral Sclerosis: Preliminary data for evidence-based guidelines. *Archives of Physical Medicine and Rehabilitation*, 91(2):268 – 272, 2010.
- [39] S. Zhang and M. Sridharan. Active visual sensing and collaboration on mobile robots using hierarchical POMDPs. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '12, pages 181–188, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.
- [40] Y. Zhang, P. Xu, T. Liu, J. Hu, R. Zhang, and D. Yao. Multiple frequencies sequential coding for SSVEP-based brain-computer interface. *PLoS ONE*, 7(3):e29519, 03 2012.