Low-Fi Skin Vision: A Case Study in Rapid Prototyping a Sensory Substitution System

Jon Bird
Pervasive Interaction Lab
Open University
Milton Keynes, UK
j.bird@open.ac.uk

Paul Marshall
Pervasive Interaction Lab
Open University
Milton Keynes, UK
p.marshall@open.ac.uk

Yvonne Rogers
Pervasive Interaction Lab
Open University
Milton Keynes, UK
y.rogers@open.ac.uk

ABSTRACT

We describe the design process we have used to develop a minimal, twenty vibration motor Tactile Vision Sensory Substitution (TVSS) system which enables blind-folded subjects to successfully track and bat a rolling ball and thereby experience 'skin vision'. We have employed a lo-fi rapid prototyping approach to build this system and argue that this methodology is particularly effective for building embedded interactive systems. We support this argument in two ways. First, by drawing on theoretical insights from robotics, a discipline that also has to deal with the challenge of building complex embedded systems that interact with their environments; second, by using the development of our TVSS as a case study: describing the series of prototypes that led to our successful design and highlighting what we learnt at each stage.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

Keywords

Lo-Fi Rapid Prototyping, Tactile Vision Sensory Substitution, TVSS, Embedded Systems

1. INTRODUCTION

Over the last forty years, Bach-y-Rita [9][10] and other researchers have developed a series of Tactile Vision Sensory Substitution (TVSS) systems that convert a camera image of a subject's environment into tactile stimulation on their bodies (Figure 1). With practice, participants are able to use this tactile information to make perceptual judgements and co-ordinate action, for example batting a ball that is rolling off a table [26]. There are three finding from TVSS research that particularly interest us. First, they vividly demonstrate that people can see what they feel: subjects consistently describe their perceptual experience in quasi visual terms, even though it is the result of somatosensory, rather than retinal,

stimulation. 'Skin vision', to use Humphrey's phrase [25], does not appear so alien when one realises that it is a phenomenon long exploited by blind people who use canes to guide themselves through real-world environments. A second interesting finding is that as participants learn to use the tactile stimulation to recognise objects or guide their behaviour, the focus of their perception shifts from their skin to surrounding space. The tactile interface becomes transparent with expert use, or 'ready-to-hand' [23]: the user experiences the world *through* the technology, which effectively becomes an extension of themselves. The third finding clearly demonstrated by sensory substitution research is that visual stimuli can be mapped to different somatosensory regions (for example, legs, back, abdomen, waist, tongue or arms) and usefully co-ordinate behaviour, as long as there is a feedback loop between the user's actions and the TVSS system. It does not matter whether the system's sensor(s) are body-mounted or positioned away from the subject: what *is* essential is that a user's own movement affects the pattern of tactile stimulation they feel.

Bach-y-Rita's pioneering research has stimulated more than his subjects' somatosensory systems. It has inspired further experiments that have established that sensory modalities other than vision can be mediated by tactile interfaces (for example, ultrasound [16] and magnetism [31]). TVSS research has also fired the imagination of philosophers (for example, [19],[32],[20]) who debate issues such as: What is the relationship between two sensory modalities that allows one to act as a substitute for the other? Given the remarkable capacity of people to adapt to changes in existing sensorimotor mappings and to incorporate novel sensory modalities, under what conditions does a sensory substitution technology *not* become 'ready-to-hand'? In the interdisciplinary e-sense project (www.esenseproject.org) our interest in TVSS systems is both practical and philosophical. One concrete goal is to build useful sensory augmentation[1] devices that are portable and can be used outside of a laboratory context. Another, more speculative, goal is to generate novel insights into the *extended mind*. This philosophical perspective argues that "the mind is less and less in the head" [19, p.6] and that human cognition emerges from a plastic hybrid of biological and non-biological components (including

---

[1]'Sensory augmentation' encompasses both 'sensory extension' and 'sensory substitution', and is where technology provides access to environmental energy not available to a person's biological perceptual system (for example, IR or ultrasound). In the substitution case this is because of perceptual impairment, for example, an individual is blind or deaf.

external representations and technologies). This viewpoint has profound implications for our notion of what it means to be human, suggesting that using new technologies can potentially change a person's thought and actions.

Figure 1: The experimental set up used by Bach-y-Rita and his colleagues in their original TVSS experiment [9]. The subjects sit in the chair and their backs are stimulated by four hundred solenoids whose activity is driven by the camera. Subjects can move the camera and also adjust its focus, thereby changing the tactile stimulation on their backs. Over time, blind subjects learn to recognise objects.

We believe that by creating a wide array of tactile interfaces and monitoring both their use and the user experiences, we can potentially gain knowledge about how to build useful sensory augmentation technologies as well as important insights into the extended mind perspective. In our interdisciplinary approach, conceptual philosophical analysis feeds into the design of systems and user studies reciprocally feed back into philosophy. Our project is extremely open-ended because: first, relatively little is known about the design issues related to portable sensory augmentation systems; and second, it is still a matter of some contention whether cognition is better understood by focusing on individual minds or on brain-artifcact systems. Clark [17] concedes that "there is a pressing need for an account of the features of an extended brain-artifact system which make it (hopefully) epistemologically sound". As a first step towards achieving both our practical and philosophical objectives, we decided to build a portable demonstration of sensory substitution. Our approach to achieving this goal is low-fi in two senses: first, we use a very low resolution vibrotactile array that contains only twenty vibration motors (commonly used in mobile phones and pagers [5]) in the latest design; second, we employ a low-fidelity rapid prototyping approach to design the technology [33].

1.1 Structure of the Paper

Having established why we are interested in sensory augmentation, in the next section we outline the requirements for our initial TVSS system. We explain why we want to build a minimal, portable system using off-the-shelf components and open source software and hardware technologies. The following section highlights some of the key technical challenges involved in designing embedded interactive systems. We then describe how the process of building these systems provides opportunities for theoretical insights. We draw on lessons that have been learnt in robotics research: a discipline where the main focus is building complex embedded systems that interact with their environments and whose insights are therefore highly relevant to our project. Our aim is to clarify the motivation underlying our employment of a rapid, low-fi prototyping methodology. In the next sections we describe the series of prototypes that led to our successful minimal TVSS system, describing the components in more detail and the results that we have achieved. We finish by describing the future directions for this work, describing a series of experiments that we are about to carry out with our system that will potentially inform philosophical questions about sensory augmentation and, more generally, the extended mind perspective.

2. REQUIREMENTS FOR OUR TVSS SYSTEM

The first requirement for our TVSS system is that it is portable. We want to be able to demonstrate sensory substitution to as many people as possible outside of a laboratory context. Froese and Spier's Enactive Torch [20] is an excellent example of a small sensory substitution device that is easily transported and it has been demonstrated to many people at conferences and workshops.

A second, more general requirement is that our TVSS system can be built by other people who only have small budgets and do not have access to specialist equipment. Our knowledge of both hardware and software has significantly increased by participating in the open source communities that developed the Processing programming language and environment [6], Arduino microcontrollers [2] and openFrameworks, a C++ library for creative coding [4]. We have benefitted from the work of many individuals who develop and maintain these tools and also from the information on techniques and applications that users post in blogs, online tutorials, videos and instructables [3]. We therefore want to contribute something back to these communities as it is the way they grow and stay vibrant. There is also the possibility that these DIY/hobbyist communities might improve our system, or transform it by using it in ways that we have not imagined. We would like to tap into the potential for distributed innovation and participatory design in these groups. Buechley [15, p.149] also advocates the 'democratisation' of design, especially as "computing moves further and further from the desktop", reminding us that "[c]omputer science
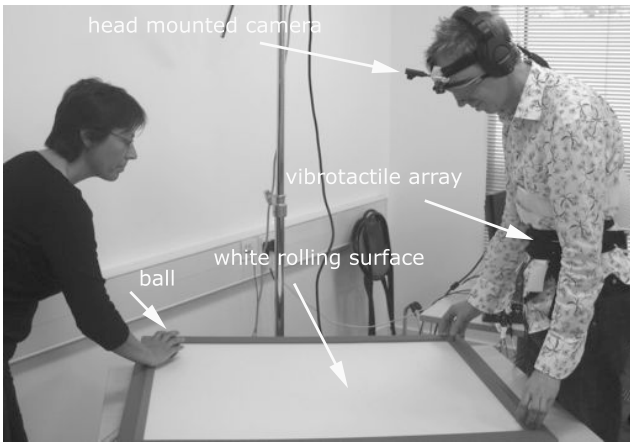
Figure 2: The first experimental set up used to test our six (2x3) vibration motor TVSS system using a two alternative forced choice paradigm. An experimenter (left) slowly rolls a black tennis ball down a gently sloping white surface towards a subject (right) wearing a vibrotactile array on their abdomen. The subject is blind-folded and wearing headphones. The activation of each motor is determined by the proportion of black pixels in one of six equal regions in the image grabbed from the head-mounted camera, each of which corresponds to a motor. The subject indicates whether they think the ball is rolling down the left or right hand side of the surface by raising the hand which they think the ball is moving towards.

has deep historical roots in the world of hobbyists, garage tinkerers, and other 'amateurs'. One can trace back innumerable important ideas in the field - in home computing, in interface design, in computational graphics - to the work of exuberantly obsessed young men and women". More specific to the design of TVSS systems, Bach-y-Rita [10, p.545] concedes that up to now sensory substitution technologies have remained at a prototype stage and urges the development of "robust and relatively inexpensive implementations of the technology to make it accessible to a wide range of patients suffering sensory loss", as well as more generally arguing that building sensory extension devices that "expand human sensibilities" will provide theoretical insights into perception. Two specific requirements for our system that are a consequence of wanting to engage open source communities are: first, it uses cheap, off-the-shelf components; and, second, any software and designs we develop are open source and freely available online.

To meet the portability and low cost requirements, we decided to keep the resolution of our system (measured by the number of vibration motors) as low as possible. In the original TVSS system, Bach-y-Rita and his colleagues used four hundred points of vibrotactile stimulation (Figure 1) [9]. In later experiments with Tongue Display Units (TDUs) the interface consisted of a 12x12 matrix of electrodes [27]. Bach-y-Rita emphasises that even "[a] poor resolution sensory substitution system can provide the information necessary for the perception of complex images. The inadequacies of the skin (e.g. poor two-point resolution) do not appear as

serious barriers to eventual high performance, because the brain extracts information from the patterns of stimulation. It is possible to recognize a face or to accomplish hand-eye coordinated tasks with only a few hundred points of stimulation" [10, p.543]. Jannson and Brabyn [26] used a 20x20 vibrotactile array in experiments where seated subjects had to bat a ball as it rolled off a table in front of them (Figure 3). The approach time was about 3.5 seconds and the two subjects achieved almost perfect performance. Importantly, although the camera position was fixed the subjects' batting hands were in its field of vision and so self-movement caused changes in the pattern of vibration they felt on their backs. In keeping with our lo-fi approach, we decided to
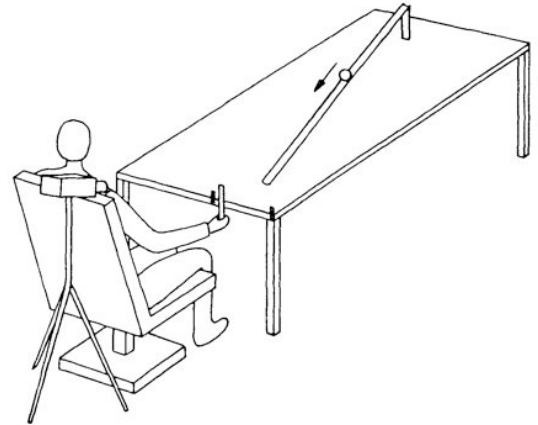
Figure 3: The experimental set up used by Jannson and Brabyn [26] in their ball batting experiment. Subjects sit on a chair and a fixed camera image is mapped onto a 20x20 matrix of solenoids that stimulate their backs. Importantly, although the camera is fixed, the subjects can detect the movement of the bat as well as the ball on their backs: the motion of their hand affects the tactile stimulation they feel. It takes about 3.5 seconds for the ball to reach the edge of the table where the subject tries to hit it with the bat.

determine the lowest resolution TVSS that would enable a subject to bat an approaching ball. As an initial task, we decided to see whether subjects could determine whether a ball was rolling down the left or right hand sides of a gently sloping surface with an extremely low resolution system, namely six points of stimulation arranged in a 2x3 matrix (Figure 4). We viewed this task as a stepping stone to developing a TVSS system that enables uses to co-ordinate faster and more complex behaviours, for example, batting. Consequently, a final requirement for our design was that we could increase the resolution of the system as necessary, while at the same time keeping its complexity and the associated cost of components as low as possible.

3. KEY CHALLENGES IN DESIGNING EMBEDDED INTERACTIVE SYSTEMS

In the next two sections, we develop an argument for adopting a low-fi rapid prototyping based on theoretical and empirical insights from robotics research. Findings from

this discipline are particularly pertinent as it has always faced the challenge of building complex embedded systems that have to interact autonomously with their environments. In this section we identify some of the key challenges in designing embedded interactive systems and argue why a prototype-driven design process is one approach to overcoming them.

Rodney Brooks famously advocated building situated robots as an alternative to good old fashioned artificial intelligence (GOFAI). He argued that a major problem with GOFAI is that the agents that are built to test theories of intelligence are essentially problem solvers that only work in abstract, symbolic domains: "The symbols may have referents in the minds of the builders of the systems but there is nothing to ground those referents in any real world. Furthermore, the agents are not situated in a world at all. Rather they are given a problem, and they solve it. Then, they are given another problem and they solve it. They are not participating in a world as would agents in the usual sense" [14, p.14]. Brooks highlighted the key challenges in designing situated agents that can participate in the real world in 'the usual sense': they have a limited perspective on the world; they have to respond quickly; and they have to deal with dynamic, noisy environments.

All three of these issues can be overlooked in high-level designs abstracted from the real world, but they became very evident when prototyping our minimal TVSS, because even though it is essentially a dumb interface for intelligent human users, it has to provide an appropriate coupling with the world otherwise it cannot be used to co-ordinate behaviour. First, as highlighted in Section 1, it is essential that users' self-motion generates changes in the vibrotactile stimulation they receive. A key design issue is deciding where to place the camera, given its limited field of vision, in order to achieve this. One solution is to head mount the camera so that a user can point it at the most relevant parts of the environment. An alternative is to use a fixed camera that sees the parts of the user's body that move during the task - following the experimental set up used by Jannson and Brabyn (Figure 3). In a later section we detail our findings from trying both these approaches. Second, the speed at which the camera images are turned into vibrations determines the subjects' response time and the maximum speed that balls can be rolled and successfully tracked and batted. Third, noisy camera images have to be pre-processed before they can be transformed into tactile stimulation that can be used to co-ordinate users' actions.

So given the design challenges identified by Brooks, what is the appropriate approach for building embedded interactive devices, such as a TVSS system, so that users can 'participate in the real world'? Hendriks-Jansen [24] highlights a further methodological issue facing robot designers: "[t]here is no formal rule for producing a particular pattern of situated activity in a particular type of environment". If we are only able to draw up high level requirements, rather than detailed design specifications, then he argues that designers have to resort to a more open-ended prototype-driven design process which is more akin to tinkering than engineering. Banzi, one of the developers of the open source Arduino microcontroller (which is a central component in all of our TVSS system designs) argues that tinkering is central to the 'Arduino Way' and provides a good description of this process: "We believe that it is essential to play

with technology, exploring different possibilities directly on hardware and software - sometimes without a very defined goal. Reusing existing technology is one of the best ways of tinkering. Getting cheap toys or old discarded equipment and hacking them to make them do something new is one of the best ways to get great results" [11].The main differences between tinkering and engineering are summarised in Table 1 (modified from [12]) and the design process we have used to develop our TVSS system is better described as the former, rather than the latter: we initially had only very general design requirements to work from, rather than a clear specification; our focus has always been to make working prototypes that we could test and refine, rather than 'best solutions'; and we use, and sometimes repurpose, readily available off-the-shelf components. Although the term 'tinkering' is used in a positive way in the open DIY/hobbyist electronics communities, the word might have negative connotations in some academic disciplines. Therefore, we also characterise our approach as 'lo-fi rapid prototyping' and in the next section we highlight how this methodology, used in software and industrial design, shares many features with tinkering.

Engineering	Tinkering
Clear goal and specification	Often no goal or specification
Aims for the best solution given constraints	Makes some kind of workable object
Not necessarily dependent on previous designs	Uses whatever is to hand
Insulates subsystems and minimises unforeseen side effects	Combines systems or transforms them for new uses

Table 1: A comparison of engineering and tinkering.

4. LOW-FI RAPID PROTOTYPING

When designing software interfaces, Rettig [33, p.23] advocates using rapidly constructed paper prototypes and testing them on users. This low-fi prototyping technique requires very few skills (cutting paper, applying glue, working a photocopier) and allows designers to test interfaces early and often: a process he describes as "a kind of natural selection for ideas". The goal is to complete as many prototype-test-refinement cycles before the design is implemented in code. The paper systems are working models, rather than illustrations, and require one tester to play the role of the computer, silently changing the paper components of the interface in response to a test user's actions. A facilitator guides the user through a task and encourages them to vocalise what they are thinking as they try out the paper prototype. On the basis of this feedback, prototype designs are refined and further testing takes place.

The low-fi materials used in the prototypes make it clear to testers that the design is 'formative' and a work-in-progress. Rettig reports that this has two advantages over 'summary evaluation' of completed products: first, users are not inhibited from suggesting changes and offering feedback, as they sometimes are when testing more polished and obviously completed systems; second, users do not just focus on 'fit and finish' issues (for example, the colours and fonts
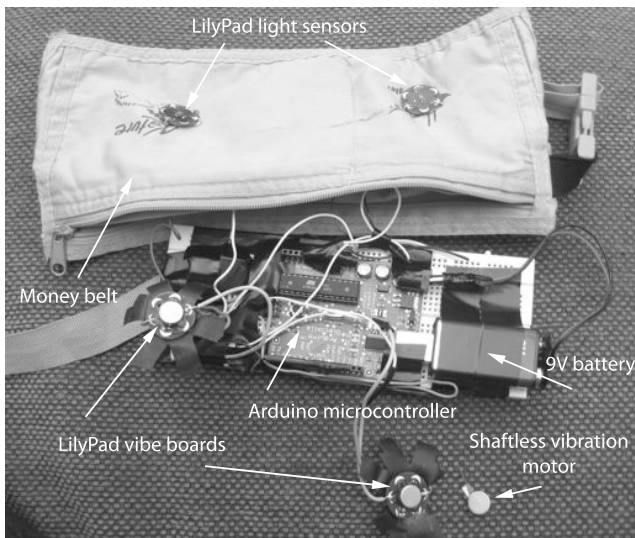
Figure 4: **Vibrating money belt: a wearable rapid prototype built to test the suitability of Arduino LilyPad vibe boards for tactile sensory extension interfaces. If light levels go above a hard-wired threshold value, then each of the sensors switches on one of the vibe boards. The diameter of the shaftless vibration motor is 10mm. The LilyPad vibe boards consist of one of these motors mounted on a printed circuit board that enables users to connect them to a microcontroller using conductive thread and incorporate them into clothing.**

used). The benefit of lo-fi paper prototyping is that the main focus is on usability issues - how effectively the software works and whether the contents of the system meet requirements. Importantly, necessary changes can be made rapidly as the systems are made out of paper. In contrast, making changes to a high-fidelity prototype can be a major challenge. Schrage [34, p.195] describes how the clay models used by American automobile companies prevented rapid iterations and modifications: "The work required to craft them made them more like untouchable works of art than malleable platforms for creative interaction". It has been speculated that this was partially responsible for the reduction in competitiveness of American car maufacturers in the 1970s and 80s. At Toyota, clay models were the *output* of CAD systems which allowed for rapid changes; in contrast in American firms the clay model was the *input* to the CAD system and became a bottle neck that prevented the effective use of computer software. Schrage argues that innovation involves shifting from what he calls 'specification-driven prototypes' to 'prototype-driven specifications'.

As part of the lo-fi prototyping approach, Rettig [33, p.24] suggests setting deadlines and sticking to them as many design issues are not apparent until testing is done with users. In order to meet deadlines decisions have to be made and if they are not the right ones then that becomes evident through user evaluation: "no matter how hard you think about it, you aren't going to start getting it right until you put something in front of actual users and start refining your idea based on their experience with your designs". Similarly,
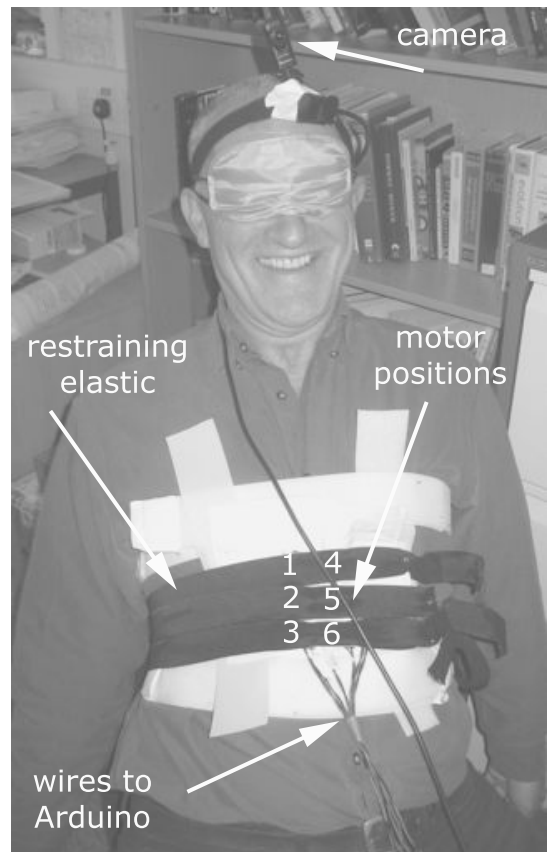
Figure 5: **An early prototype minimal TVSS system. The vibrotactile array consists of six shaftless vibration motors, each of which is controlled by a separate Arduino microcontroller PWM channel. The motors are fixed to cloth with gaffer tape and the array is held on the subject's abdomen using velcro straps. Three elastic straps keep the motors pressed against the subject. The activity of the motors is determined by the image captured by the head-mounted camera.**

Schrage [34] outlines the advantages of 'periodic prototyping' where there are regular dates when prototypes have to be delivered. This increases the number of prototypes built as well as providing a way of measuring progress. The pressure to produce some sort of workable prototype encourages a shift to prototype-driven specifications, which is generally a more open-ended and exploratory approach to design. We embrace periodic prototyping, aiming to produce a new prototype every week and demo it at our weekly informal meeting where all researchers in the lab meet for coffee and a 'show and tell' session. This generates a creative pressure that drives the prototyping process. Importantly, the regularity of these meetings means that we are not expected to demonstrate major developments each time: even small changes to the system are appropriate for this forum where many projects are being discussed.

To summarise, low-fi prototyping is an effective methodology because: it puts a focus on usability issues; it increases the number of prototype-test-refinement cycles; and designs

are seen as formative or in-progress rather than finished and immutable. Regular prototyping deadlines increase the number of test cycles and encourage an exploratory approach. A rapid prototyping/tinkering approach is particularly effective for building embedded systems that are situated in, and interact with, the real-world. This is because there are rarely clear specifications for this sort of system and a designer must respond to the difficult challenges identified by Brooks (the systems have a limited perspective on the world; they must respond rapidly; and the world is noisy and dynamic). Arguably, rapid prototyping may be the only viable way to proceed when designing this sort of system. However, this design approach is not just a last resort: we also argue that it provides opportunities for developing novel theoretical and philosophical ideas. This is of particular relevance to our project and is the focus of the next section.

5. ADVANTAGES OF BUILDING PHYSICAL PROTOTYPES

A fundamental advantage of building physical devices and prototyping them in the real world is that this facilitates the modelling of agent-environment interactions [35]. For example, the primary reason that robots can be fruitful sources of scientific hypotheses in the biological sciences is that they are embedded in the physics of the real world and can give us insights into what different environments afford organisms. Webb [36, p.5] emphasises that it is not always obvious which aspects of a phenomenon to abstract into a model and that in conventional science "the purpose in modelling is often to discover what are the 'relevant features' or 'essential structures', so model usage cannot depend on prior knowledge of what they are to establish the modelling relationship". Similarly, McIver [28, p.1067] highlights how recent biorobotics experiments that approximate the movement of fish tails and insect wings in fluids have facilitated the abstraction process and led to the discovery of fluid phenomena. He argues that by "allowing the full complexity of the environment to work on what could be called 'reduced robotic preparations', this research is cracking open the black box of complex deformable-body and fluid dynamics phenomena to new theoretical advances. The epistemic accessibility afforded by building these robotic devices is analogous to that obtained by traditional instruments such as the microscope and telescope. The building of physical models not only reduces abstraction load, but in illuminating that part of nature we most urgently need to abstract in order to account for a phenomenon, it provides a saliency filter for the immense richness of opportunities for abstraction effort that arise at every turn in the course of experimental work".

5.1 Philosophy using a Screwdriver

Building embedded interactive systems can facilitate the development of theoretical frameworks, such as the extended mind. For example, in the field of Evolutionary Robotics (ER), Harvey advocates building robots to gain insights into cognition: what he calls 'Philosophy using a screwdriver' [21]. He and his colleagues argue that physical robot models of cognition, designed using genetic algorithms, can provide existence proofs that demonstrate the sufficient conditions for phenomena to occur. They argue that these existence proofs can then be used for "catalysing theoret-

ical re-conceptions and facilitating the production of novel hypotheses which then need to be appropriately translated to domain-specific cases and tested empirically" [22, p.7]. Similarly, in our project we believe that prototyping sensory augmentation systems can help to address what Clark [18, p.29] has identified as the most important issue in extended mind research, namely to identify and understand, "the range and variety of types of cognitive scaffolding, and the different ways in which non-biological scaffolding can augment (or impair) performance". Specifically, we would like to understand the conditions under which sensory augmentation technologies become transparent.

Menary, an advocate of the extended mind, argues that one reason that there is some resistance to this philosophical perspective is that "a brain-tool science would have to cover a motley of processes and thus could not be the basis of genuine scientific enterprise" [29, p.338]. One way to characterise this problem is that there is no 'natural kind' that seems to capture the diversity of brain-artifact systems. Nagel [30, p.31] suggests how building systems can address this issue when he argues that "[t]he development of comprehensive theoretical systems seems to be possible only after a preliminary classification of kinds has been achieved, and the history of science repeatedly confirms the view that the noting and mutual ordering of various kinds - a stage of enquiry often called 'natural history' - is a prerequisite for the discovery of more commonly recognized types of laws and for the construction of far-reaching theories". To summarise, building sensory augmentation technologies and testing them in the real world facilitates an understanding of agent-environment interactions by reducing 'abstraction load' and potentially helping discover relevant features. Furthermore, this approach can help develop and test theoretical assumptions.

6. THE EVOLUTION OF OUR MINIMAL TVSS SYSTEM

Having argued for employing a low-fi rapid prototyping approach to designing embedded interactive systems, in the following sections we now describe how we have applied this methodology to building a minimal TVSS system. We describe the series of prototypes we have built and how they have informed succeeding designs. Our current design (Figure 6) meets the requirements detailed in Section 2 and was successfully demonstrated at Brighton Science Festival on 21 February, 2009, where 100 primary school children (and a few of their parents) wore the system and experienced 'tummy vision'[2] over an 8 hour period (Figure 7).

6.1 Prototype 1: Vibrating Moneybelt

This prototype (Figure 4) was built to test the suitability of Arduino LilyPad [1] vibe boards for building wearable vibrotactile interfaces. The LilyPad vibe boards consist of a 10mm shaftless DC motor mounted on a thin printed circuit board that can be connected to a microcontroller using conductive thread and incorporated into clothing. In our simple device, if either of the photo sensors detects light levels

---

[2]We advertised our demonstration with the question 'Can you see with your tummy?', a word more appropriate than 'abdomen' for our target audience of 7-11 year olds. It clearly worked as there was more than a two hour waiting list to try out our TVSS system.

above a hard-wired threshold value, then an Arduino micro-controller [2] switches on the vibe board they are connected to. Inspired by Braitenberg's classic *Vehicles* [13], this system can act as a light seeking (vehicle 2a) or light avoiding (2b) device, depending on whether the light sensors connect to the motor on the same or opposite side of the money belt. The vibration was very clear through clothes. However, even though the vibe boards are relatively inexpensive components (£10 each), we decided to buy the vibration motors directly (£2.50 each) in order to save money and enable us to potentially build large vibrotactile arrays.

6.2 Prototype 2: PWM Controlled Mobile Phone Vibrators

We attached a 10mm shaftless DC motor, commonly used in mobile phones [5], to perf board, along with a snubber diode to protect the microcontroller. These had one connection to a 3.8V power supply and another to a transistor which was driven by an Arduino pulse width modulation (PWM) channel. By varying the PWM signal it was possible to control the intensity of vibration, although frequency and amplitude cannot be separately adjusted. We attached these motors to various body areas (head, abdomen, back and arms) using tape or simply holding them in place by hand. We found that although there was variation in individual sensitivity, subjects (first author, his children and various lab members) were able to discriminate between three and five intensities (essentially, low, medium and high).

6.3 Prototype 3: Vibrotactile Belt

We attached six vibrators to a strip of elastic, held in place with a safety pin, and controlled all of them via separate PWM channels on an Arduino microcontroller. This design was adjustable so that it could be worn by both children and adults of varying sizes. Even through clothes, subjects could discriminate low, medium and high intensities. They could easily detect movement around the belt but were not always able to localise the stimulation to a particular vibrator.

6.4 Prototype 4: Head-Mounted Webcam

We chose a Soyntec Joinsee 110 webcam [7] as it is small and light and has a flexible stem that allows it to be easily positioned. Using gaffer tape we attached it to an elastic head torch harness. We used a 10m USB extension cable to connect the camera to a laptop. We tested two candidate image processing algorithms that we thought would detect the motion of a rolling ball using Processing [6]. A rapid comparison was possible because we had existing code that grabbed a camera frame and the Processing website has example computer vision code online. By observing the results in real-time, it was quickly evident that frame differencing is a good technique if the camera is static, but as soon as camera can move, the results get very noisy because of self-produced motion. Rather than increasing software complexity and trying to measure self-motion and compensate for it, we instead simplified the environment and used thresholding. When a dark ball is moving on a white surface it is easy to track its movement in a thresholded image. If a user positions the camera over the rolling surface then most of the image is white and if they make small head movements the only thing that changes is the position of the ball and the location of the associated stimulation in the vibrotactile array.

6.5 Prototype 5: Stimulating a 2x3 Vibrotactile Array

The head-mounted webcam captures a 321 x 240 pixel image which is then split into 6 equally sized cells (107 pixels high by 120 pixels wide) which are thresholded and analysed. The thresholding algorithm compares the grey level of each pixel to a threshold value. If the pixel is below the threshold, it is set to black, otherwise it is set to white. If a cell contains more than a specified percentage of black pixels, the corresponding vibration motor in the vibrotactile array is activated (Figure 5).

In a first experiment, an 8 year old female blind-folded subject (the first author's daughter) wore the TVSS and she had to indicate by raising her hand whether a black tennis ball was rolling down the left or right hand side of a white, gently inclined surface. It was found that there had to be sufficient light in the room, otherwise the thresholded image was too noisy. Consequently, a change was made to the software to make the threshold value adjustable. Initially, the balls were rolled too quickly and the subject was unable to track the movement because of the lag between the image capture and activation of the vibrotactile array. It was also found that successful tracking of the balls was primarily due to listening to their movement, rather than using the TVSS system, and so in subsequent experiments subjects wore headphones.

After some fine tuning of the image processing software, a second experiment was conducted with a blind-folded adult male subject who wore headphones (Figure 2). He stood at the end of a table and a black tennis ball was rolled down a very gently sloping ramp. We used the same two alternative forced-choice procedure adopted in experiment one: the subject had to raise the hand corresponding to the side of the ramp that the ball was moving down. The subject positioned his hands so that they did not touch the ramp in case he could pick up vibrations from the rolling ball. He did 25 trials, with the ball randomly rolled down the left or right side of the ramp. In 19 of the trials (76%) he correctly indicated which side of the ramp the ball was rolling down before it reached the end of the ramp. In the unsuccessful trials the subject made no indication at all, rather than falsely indicating where the ball was rolling. In one case, this was because the ball rolled down the middle of the ramp, activating both the left and right hand side motors. In the other cases the subject had not positioned his head appropriately and so he was receiving a lot of activation on the vibrotactile array from dark objects in the test environment. This suggested that if subjects could indicate when they were ready, many could rapidly achieve a perfect performance on this forced-choice task.

There were three main findings. First, the initial position of the camera is critical if the ball is to be detected: if it is not over the white rolling surface then there is too much noise in the image. Finding this 'quiet spot' was a major challenge for the subject. Second, the 2x3 resolution was not sufficient to discriminate balls rolling down the middle of the surface. Third, the lag between image capture and vibration motor activation meant that subjects could only determine which side the balls were rolling down when they were close to the edge of the table. Unsurprisingly, at no point did the subject feel as though their focus had shifted from the TVSS interface to the environment at large: the TVSS system never became transparent.

6.6 Prototype 6: Faster Image Processing

Even with fine tuning, the Processing computer vision code was slow and meant that only slowly moving balls could be tracked. We therefore re-wrote the program in open-Frameworks [4], an open source C++ library which also incorporates openCV, an open source computer vision library created by Intel. This code runs significantly faster than Processing. Our next step was to conduct experiments with this software driving the TVSS system.

6.7 Prototype 7: Colour Tracking

One fundamental challenge is what Jannson and Brabyn [26] call the 'masking problem', that is, the way that the perception of an object or target can be hidden by other environmental features. In our second, successful ball tracking experiment, all but one of the failures were due to the subject not being able to initially position his head so that the thresholded camera image was predominantly white and no stimulation was sent to the vibrotactile array. If the subject was able to find this 'still' position before the ball was released, then they were able to track the ball without any problem. However, because of the difficulty of pointing the head-mounted camera, users remained focused on the interface and the TVSS system never became transparent. We therefore used the openCV library, developed by Intel, to implement colour tracking of the ball and the user's hand. With this computer vision technique we can also use a fixed camera while still allowing a user's self-motion to affect the vibrotactile stimulation they receive by tracking the motion of one of their gloved hands. We found that a bright orange juggling ball and a fluorescent yellow cycling glove could be effectively tracked in a range of environments and lighting conditions. The software identifies the centre of the ball and the glove and activates the vibration motor in the array which corresponds to this point. We mapped a strong vibration to the ball and a weaker one to the hand and most subjects were able to discriminate the ball movement from their self-motion. Some subjects found it hard to sense their hand movements if they were wearing thick clothing.

6.8 Prototype 8: Larger Vibrotactile Arrays Using a TLC5940 chip

In order to overcome the limitation of only 6 PWM channels on the Arduino, in this prototype we used multiplexer chips to increase the number of channels. A library has been developed by Alex Leone, a member of the Arduino community, [8] that makes it relatively simple to control the Texas Instruments TLC5940 chip which was designed for controlling LED arrays. Each chip has 16 channels and multiple chips can be daisy chained together. Some Arduino users, adopting a tinkering approach, have repurposed this chip to control DC vibration motors and we successfully controlled 16 vibrators using one chip.

6.9 Prototype 9: Custom PCBs

With the help of Bill Bigge (Autonomous Systems Lab, University of Sussex) we designed custom printed circuit boards (PCBs) to miniaturise the circuits we developed in earlier prototypes so they could be more easily incorporated into a wearable system. We also developed a circuit board with surface mounted components that reduces the size of a vibration motor unit. Bill built an Arduino shield that allows multiple TLC5940 chips (up to 16) to be daisy chained

and controlled from a single microcontroller. In this prototype we wired up 16 vibration motors to determine how close we could place them on the abdomen and the wearer still distinguish individual points of stimulation. One of the biggest challenges was keeping the 'wiring spaghetti' under control, but the PCBs have built in connection headers to help minimise this issue.
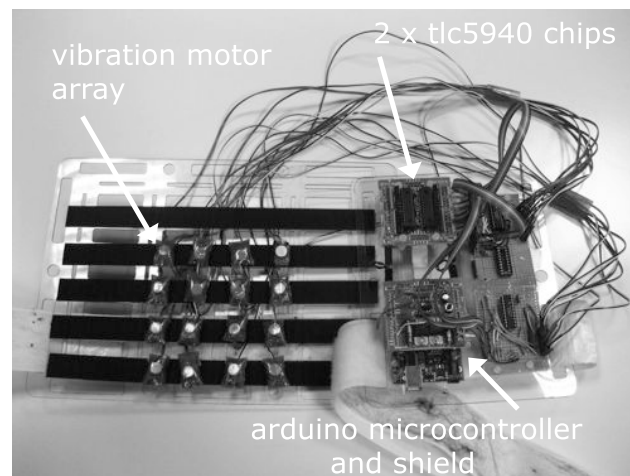
Figure 6: Prototype 10: the minimal TVSS system that was successfully demonstrated on over 100 children and adults at the Brighton Science Festival, 2009. The motors are attached to a piece of polycarbonate by velcro, allowing rapid reconfiguration of their pattern and number. This is held against the users abdomen by an adjustable velcro belt. The Arduino microcontroller, battery and other electronics are also attached to this belt.

## 6.10 Prototype 10: Remodelled 5x4 Vibrotactile array

We made the structures that hold the electronics and vibrotactile array more robust: replacing gaffer tape and felt with laser cut polycarbonate (Figure 6). The motors are now attached using velcro which makes it easy to reconfigure their pattern and number.

We discovered that low current motors can be run directly from the TLC5940s if sufficient diodes are included in the circuits, whereas before we had been decoupling the chips from the motors using transistors or logic buffers. This reduces the cost and complexity of the system.

This prototype was thoroughly tested on Saturday 21 February, 2009, when we ran a 'tummy vision' demonstration at the Brighton Science Festival. Over 100 children (and a few parents) tried to bat a ball rolling towards them with their hand, a task similar to that used by Jannson and Brabyn [26] (Figure 3). We used a fixed camera to reduce the set up time for each user. We built a rolling ramp to control the speed of the balls (Figure 7). Initially, we tested each user's ability to bat a ball with their eyes open. We primed them by saying '3, 2, 1, go' before we released the ball. We then blindfolded the user to demonstrate how this impaired their ability to bat the ball. We instructed people to try to bat the ball only when they thought they knew where it

was and not to use random swiping motions. Finally, people put on the TVSS system and after a few trials, nearly all of them were able to hit the ball on some trials and experience sensory substitution. We noticed that there were some individual differences in batting ability and generally children were better than adults, perhaps because of a lack of self consciousness in the demonstration environment. The children were highly motivated to do well, treating the demo as a game (and occasionally trying to cheat by moving their head back in an attempt to see under the blindfold). The stand was hugely popular and the waiting list was over 2 hours but everyone seemed to enjoy the novel perceptual experience. For 8 hours, the minimal TVSS worked effectively in a wide range of light conditions and on very different body sizes. We did not have to change the 9V battery powering the system. The only issue we found was that occasionally one of the plugs connecting the motors to the TLC5940 chips would come loose when we were removing the system from a participant. This will be simple to rectify in a future prototype.
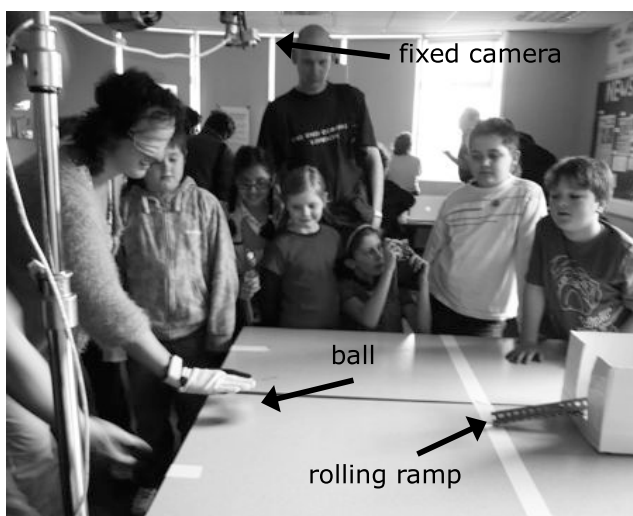
Figure 7: An adult trying out the minimal TVSS system at the Brighton Science Festival 2009. The camera is fixed overhead and the motion of the subject's gloved hand and the ball are mapped to stimulation on a 5x4 array of vibration motors worn on the abdomen. Even with this low resolution, people rapidly learn to bat the ball as it rolls towards them.

The festival clearly demonstrated that we have achieved the requirements described in Section 2: we have built a portable, low cost TVSS system from readily available components using open source software and hardware. A 5x4 array is sufficiently large to enable most people to quickly learn how to track and bat a relatively slow moving ball. The system works with both a head-mounted and fixed camera. Although a head-mounted camera is more portable (it doesn't require a stand), most users find the fixed camera easier to use as the head-mounted system requires some experimentation to find the most effective head and body position to track the ball. In a future prototype we will explore whether head-mounting a wider angle camera will improve usability.

7. FUTURE WORK

We are going to build prototypes that explore how to incorporate batteries and wiring as unobtrusively as possible into a wearable system. One approach we are actively exploring is building our own thin, flexible printed circuit boards that will contain the power and control circuitry for small groups of motors. From these modules, larger vibrotactile arrays could be constructed. We plan to investigate whether users can track significantly faster balls with larger arrays.

Given the robustness of prototype 10, we can now use the system to conduct novel experiments into sensory substitution that will potentially inform the philosophical questions motivating our research. As mentioned in Section 1, a key issue we want to explore is the conditions under which the system becomes transparent in use. This will require people to use the system for extended periods of time. The Brighton Science Festival demo illustrated the value of a game-like task for motivating users and we plan to build another minimal TVSS so that two users can play a ball batting game similar to Air Hockey. Following Bach-y-Rita, we plan to identify when the technology becomes transparent by getting people to report their experience of using the TVSS system. In the transparent case their focus will be on the environment; in the opaque case their focus will be on the interface itself.

We intend to manipulate the mapping between the camera and the vibrotactile array in software so that we can investigate how transparency and batting performance are correlated. For example, by rotating the mapping by 90 degrees clockwise the stimulation that was at the top of the abdomen will then be positioned on the right, and the stimulation that was on the left will be positioned at the top. We predict that this will disrupt transparency and reduce batting performance. We hypothesise that this change in the mapping will be adapted to more quickly when the user is able to affect the stimulation they receive by self-motion. We plan to investigate whether self-motion generated by a head-mounted camera and that generated by tracking a user's hand are equally effective at enabling a user to adapt to mapping changes. These experiments are simple to carry out using our TVSS system and will provide data that can inform the extended mind framework.

8. CONCLUSION

We have described the technical and philosophical motivations underlying our investigations of sensory augmentation technologies. Adopting a low-fi, rapid prototyping approach has enabled us to construct a minimal TVSS in five months that met our initial requirements: it is portable, low cost and constructed from off-the-shelf components using open source software and hardware. This system has enabled a large number of users, ranging in age from 6 to over 60, to experience sensory substitution and perform a ball batting task in a public context. Our next steps are to: increase the size of the vibrotactile array; develop another system so that two users can engage in competitive ball batting games: and explore the conditions under which the system becomes transparent in use. We believe our minimal TVSS system demonstrates that a low-fi rapid prototyping design methodology can overcome the challenges involved in building embedded interactive systems as well as potentially pro-

vide opportunities for gaining theoretical insights into the extended mind.

9. ACKNOWLEDGEMENTS

10. REFERENCES

[1] Arduino lilypad sewable electronic components http://web.media.mit.edu/ leah/lilypad/.

[2] Arduino open source microcontrollers http://www.arduino.cc.

[3] Instructables user-created and uploaded do-it-yourself projects http://www.instructables.com.

[4] openframeworks open source c++ library for creative coding http://wiki.openframeworks.cc/.

[5] Precision microdrives - supplier of vibration and other motors http://www.precisionmicrodrives.com.

[6] Processing open source programming language and environment http://www.processing.org.

[7] Soyntec joinsee 110 webcam http://www.soyntec.com.

[8] Tlc5940 arduino library http://code.google.com/p/tlc5940arduino/.

[9] P. Bach-y Rita, C. C. Collins, F. A. Saunders, B. White, and L. Scadden. Vision substitution by tactile image projection. *Nature*, 221:963 – 964, 1969.

[10] P. Bach-y Rita and S. Kercel. Sensory substitution and the human machine interface. *Trends in Cognitive Sciences*, 12(12):541 – 546, 2003.

[11] M. Banzi. *Getting Started with Arduino*. O'Reilly, 2009.

[12] J. Bird and P. Layzell. The evolved radio and its implications for modelling the evolution of novel sensors. In D. Fogel, M. El-Sharkawi, X. Yao, G. Greenwood, H. Iba, P. Marrow, and M. Shackleton, editors, *Proceedings of the 2002 Congress on Evolutionary Computation (CEC 2002)*, pages 1836–1841, Piscataway, NJ, 2002. IEEE Press.

[13] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. The MIT Press, Cambridge, MA, 1987.

[14] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139 – 159, 1991.

[15] L. Buechley and M. Eisenberg. Fabric pcbs, electronic sequins, and socket buttons: techniques for e-textile craft. *Personal and Ubiquitous Computing*, 13:133–150, 2009.

[16] A. Cassinelli, C. Reynolds, and M. Ishikawa. Augmenting spatial awareness with haptic radar. In *Tenth International Symposium on wearable computers*, pages 61 – 64, 2006.

[17] A. Clark. Reasons, robots and the extended mind. *Mind and Language*, 1(1):21–33, 2001.

[18] A. Clark. Towards a science of the bio-technical. *Cognition and Technology*, 1(1):21–33, 2002.

[19] A. Clark. *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford University Press, Oxford, 2003.

[20] T. Froese and A. Spiers. Toward a phenomenological pragmatics of enactive perception. Cognitive Science Research Paper 593, University of Sussex, 2007.

[21] I. Harvey. Robotics: Philosophy with a screwdriver. In T. T. Gomi, editor, *Evolutionary Robotics: From Intelligent Robots to Artificial Life. Proceedings of the 8th International Symposium on Evolutionary Robotics (ER 2001)*, volume III, Toronto, Canada, 2000. AAI Books.

[22] I. Harvey, E. Di Paolo, E. Tuci, R. Wood, and M. Quinn. Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1-2):79–98, 2005.

[23] M. Heidegger. *Being and Time*. Harper and Row, 1962.

[24] H. Hendriks-Jansen. *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought*. The MIT Press, Cambridge, MA, 1996.

[25] N. Humphrey. *A history of the mind*. Simon and Schuster, New York, 1992.

[26] G. Jannson and L. Brabyn. Tactually guided batting. Uppsala Psychological Reports 304, University of Uppsala, 1981.

[27] S. Kercel and P. Bach-y Rita. Non-invasive coupling of electronically generated data into the human nervous system. In *Wiley Encyclopedia of Biomedical Engineering*. Wiley, 2003.

[28] M. MacIver. How building physical models can reduce and guide the abstraction of nature. *Behavioral and Brain Sciences*, 24(6):1066–1067, 2001.

[29] R. Menary. Attacking the bounds of cognition. *Philosophical Psychology*, 19(3):329–344, 2006.

[30] E. Nagel. *The Structure of Science: Problems in the Logic of Scientific Explanation*. Routledge, London, 1961.

[31] S. K. Nagel, C. Carl, T. Kringe, R. Martin, and P. Konig. Beyond sensory substitution - learning the sixth sense. *Journal of Neural Engineering*, 2:13–26, 2005.

[32] A. Noe. *Action in Perception*. MIT Press, Cambridge, MA, 2004.

[33] M. Rettig. Prototyping for tiny fingers. *Communications of the ACM*, 37(4):21–27, 1994.

[34] M. Schrage. Bringing design to software. In *Cultures of Prototyping*. ACM Press, 1996.

[35] T. Smithers. Are autonomous agents information processing systems? In L. Steels and R. Brooks, editors, *The Artificial Life Route to Artificial Intelligence*, pages 123–162. Lawrence Erlbaum, Hillsdale, NJ, 1995.

[36] B. Webb. Can robots make good models of biological behaviour? *Behavioiural and Brain Sciences*, 24(6):1033–1050, 2001.