

Sensory Substitution and Perceptual Learning

By Kevin Connolly

Abstract: When a user integrates a sensory substitution device into her life, the process involves perceptual learning, that is, ‘relatively long-lasting changes to an organism’s perceptual system that improve its ability to respond to its environment’ (Goldstone 1998: 585). In this paper, I explore ways in which the extensive literature on perceptual learning can be applied to help improve sensory substitution devices. I then use these findings to answer a philosophical question. Much of the philosophical debate surrounding sensory substitution devices concerns what happens after perceptual learning occurs. In particular, should the resultant perceptual experience be classified in the substituted modality (as vision), in the substituting modality (as auditory or tactile), or in a new sense modality? I propose a novel empirical test to help resolve this philosophical debate.

Keywords: Sensory Substitution; Perceptual Learning; Learning; Training; Attention; Sensory Substitution Devices; SSDs; Integration; Attentional Weighting; Latent Inhibition

1. Introduction

Sensory substitution devices [SSDs] deliver information about the environment by converting the information normally received through one sense modality into information for another sense modality. Typically SSDs are used by the blind to help them better navigate the world. Such devices consist of a video camera which feeds visual information into a converter that transforms it into auditory or tactile stimuli, which the blind subject can perceive. In such a case, audition or touch acts as the ‘substituting modality’ for vision, the ‘substituted modality.’

When a user integrates an SSD into her life, the process involves perceptual learning: ‘relatively long-lasting changes to an organism’s perceptual system that improve its ability to respond to its environment’ (Goldstone 1998: 585). Integrating an SSD involves long-lasting changes to a user’s perceptual system (see, for instance, research on sensory substitution and brain plasticity, especially Ptito et al. 2005, and Kupers and Ptito 2004). Furthermore, these long-lasting changes allow the user of a sensory substitution device to better respond to stimuli in her

environment. A blind person, for instance, can use an SSD to respond to the colour of objects, where without the device she could not.

In this paper, I pose a challenge for sensory substitution researchers, and offer a positive proposal for how researchers might respond to the challenge. The challenge is that to understand and improve the kind of perceptual learning required for sensory substitution, researchers need to focus as much on the *unlearning* required as on the learning. That is, they should start by recognizing the long-lasting changes that have already occurred in a subject's perceptual system prior to her first experience with an SSD. I detail three such long-lasting changes in section two: latent inhibition (2.1), attentional weighting (2.2), and unitization/differentiation (2.3). I argue that these three perceptual habits are not adapted for sensory substitution and are difficult to break. My positive proposal, however, is that researchers might circumvent these habits by using SSDs that employ novel stimuli, and I outline an empirical way to determine which stimuli are novel. Researchers might speed up the perceptual learning process by reducing the amount of unlearning that sensory substitution requires.

The argument in section two is in service of a philosophical point, which I make in section three. Much of the current philosophical debate surrounding SSDs is about the nature of one's perceptual experience once one has fully integrated an SSD. The question is whether one's perceptual experience should be classified in the substituted modality (as vision), in the substituting modality (as auditory or tactile), or in a new modality (see Hurley and Noë 2003; Block 2003; Noë 2004; O'Regan 2011; Kiverstein, Farina, and Clark forthcoming; and Deroy and Auvray forthcoming). I argue that we can use an empirical test for determining novel stimuli to make progress towards resolving this philosophical debate.

2. Perceptual Learning Effects

Perceptual learning involves long-lasting changes to one's perceptual system. In this section, I identify three specific kinds of long-lasting changes that will have already occurred in a subject's perceptual system prior to her first experience with an SSD: effects of latent inhibition, attentional weighting, and unitization/differentiation. I show some ways in which each of these effects would need to be undone for sensory substitution learning, but I also show how researchers can circumvent the need to undo them if they program SSDs such that the devices deliver novel stimuli.

2.1 Latent Inhibition

A prominent model of perceptual learning is Ian McLaren and Nick Mackintosh's associative learning model (see McLaren, Kaye, and Mackintosh 1989, and McLaren and Mackintosh 2000). McLaren and Mackintosh explain how perceptual learning occurs in cases such as learning to discriminate a Merlot from a Cabernet Franc. Suppose you are presented with two similar perceptual stimuli (X and Y) the same number of times. Being similar stimuli, many features are shared between X and Y. When a feature is shared, you are exposed to it twice as much as you are a non-shared feature (once when presented with X, and a second time when presented with Y). According to the McLaren and Mackintosh model, you become habituated to these shared features, which results in those features becoming less likely to enter into associations (and as such, less likely to be good candidates for use in learning). On the other hand, you are exposed to the non-shared features only half as much as the shared features. So, you become less habituated to the non-shared features. This results in the non-shared features becoming more likely to enter into associations (and so more likely to be good candidates for use in learning).

To take a concrete example, suppose that you drink a Merlot and a Cabernet Franc the same number of times. The features shared in common between them, such as the tannins, are exposed twice as much as the features unique to each of them. You become more habituated to these shared features than to the unique ones, resulting in the shared features becoming less likely to enter into associations. As such, the shared features are poor candidates for use in learning. In contrast, you become less habituated to the features unique to each of the wines. So the unique features, such as the fruitiness of the Merlot, are more likely to enter into associations. As such, the unique features are better candidates for use in learning.

McLaren and Mackintosh's model is strikingly counter-intuitive, despite being well-supported by experimental evidence. On their view, it is easier to learn from unfamiliar stimuli and harder to learn from familiar stimuli. It seems at first glance that it would be the other way around. And, in fact, for several decades in the history of psychology, that intuitive idea was dogma.

In the late 1950s, psychologists were concerned with the question: Can learning take place without reinforcement?¹ There was evidence for this. One important experiment was conducted by E.C. Tolman and H.C. Honzik in 1930. The experiment involved rats and mazes, and had two stages. In the first stage, rats were allowed to explore a maze without reinforcement (that is, there was no reward at the end of the maze). In the second stage, the rats were rewarded with food at the end of the maze. These rats were more quickly able to navigate the maze and receive the reward than a control group that did not take part in the first stage. The idea was that the rats had learned about the routes of the maze during the first stage, even without a reward, and that learning became evident only in the second stage when a reward was offered. These rats

¹ My exposition in this section is based largely on Lubow and Weiner, 2010.

experienced what was called ‘latent learning,’ a kind of learning that occurs before reinforcement, but is evident only after reinforcement (Tolman and Honzik 1930: 257).

In 1959, R.E. Lubow and A.U. Moore were attempting a simple demonstration of latent learning, not with rats, but with sheep and goats (see Lubow and Moore 1959). They divided the animals into two groups. In the first stage of the experiment, they presented each member of the first group with a flashing light ten times, and presented each member of the second group with a rotor that turned ten times. None of the ten presentations were reinforced for either group. In the second stage of the experiment, both groups received alternated presentations of the flashing light and the turning rotor. For each presentation, the flash or the turn was followed by a mild shock.

Lubow and Moore were expecting their results to provide an example of latent learning. Specifically, they thought that it would be easier to condition the animals with the stimulus that they were pre-exposed to in stage one—the familiar stimulus. In fact, they found exactly the opposite. Their original hypothesis was that just as the rats in Tolman and Honzik’s experiment had learned to navigate the maze even without reinforcement (something that became evident only after reinforcement in stage two), so too would the sheep and goats learn about the stimuli they had been exposed to in stage one, and this learning would become evident after the reinforcement in stage two. So, for instance, if the animal had been presented with a flashing light ten times in stage one, Lubow and Moore predicted that in stage two, the animal would build an association between the flashing light and the shock more quickly than between the turning rotor and the shock.

What Lubow and Moore actually found was the inverse of what they had predicted. If the animal was pre-exposed to the flashing light, then it built an association between the turning

rotor and the shock more quickly. If it was pre-exposed to the turning rotor, then it built an association between the flashing light and the shock more quickly. To contrast with latent learning, Lubow and Moore called the effect that they had found ‘latent inhibition.’ In latent inhibition, pre-exposure to a stimulus type inhibits learning on that stimulus type.²

Latent inhibition supports McLaren and Mackintosh’s model of perceptual learning. Their model makes the prediction that it is harder to learn from familiar stimuli. As the story of Lubow and Moore’s experiment indicates, latent inhibition is a counter-intuitive phenomenon, but it is also a well-established one. While it seems like familiar stimuli would be easier to learn from familiar stimuli, in fact the opposite is true.

2.1.1 Latent Inhibition and Sensory Substitution

a. Harnessing Latent Inhibition for More Efficient Learning with Sensory Substitution Devices

What can sensory substitution researchers learn from studies on latent inhibition? One lesson is that effective perceptual learning is counter-intuitive, and so the way to most effectively train subjects on SSDs may defy common sense. In particular, since perceptual learning is quicker on novel stimuli due to latent inhibition, researchers might increase the speed of the learning process with SSDs by programming the most novel stimuli into SSDs. To do this, researchers need to be able to test that a stimulus programmed into an SSD is novel to subjects, as opposed to familiar. I want to suggest two ways to do this.

One way to determine whether stimuli are novel to subjects is by examining their brain responses. Novel stimuli activate networks in the brain that have a signature and are detectable by using electrophysiological recordings of scalp and intracranial ‘event-related potentials’ (ERPs). These signatures have been specifically shown to exist for novel auditory and tactile

² Fifty years later, Lubow is surprised that the original experiment worked at all, given that the number of stimulus pre-exposures was so small (2010, p. 4). Typically latent inhibition requires many more than ten exposures.

stimuli (Knight 1996). ERPs are not the only option, however. For cortically blind subjects at least, another option might be to examine pupil response to determine novel stimuli. Naber and colleagues have provided evidence that the pupil dilates more for novel stimuli (Naber et al. 2013), and there is evidence that pupil dilation measures novelty not just for visual stimuli, but for other modalities as well (see Naber et al. 2013: 2). This could provide a helpful test for sensory substitution researchers. As Naber and colleagues point out, pupil size is ‘a simple and outwardly accessible physiological measure’ (2013: 2). By measuring pupil dilation or by using ERPs, researchers might determine which stimuli are novel to subjects, and program SSDs such that they deliver those stimuli. (Note that this will be important at the end of the paper, when we engage with the philosophical debate over whether SSD perception should be classified as visual, as auditory or tactile, or as part of a new modality.)

While I think ERPs or pupil dilation measurements are the ultimate standard as to which stimuli are novel and which are not, even without that data, we can make some progress on the question. In ordinary perceptual development, for instance, we can say that the subject learns on novel stimuli. The world of an infant is filled with shapes, sounds, colours, tastes, textures, and smells, and for each of those properties, there is a time that they are novel to her. The first experience of green is new, as is the first feel of a block, and the first perception of the sound of a barking dog. Contrast this with sensory substitution. In sensory substitution, the subject learns on proximal stimuli³ to which, in some sense, he has already been exposed.

Consider the vOICE, ‘a visual-auditory substitution device that works by transforming images from a digital camera embedded in a pair of sunglasses into auditory frequencies

³ I take a stimulus to be an objective property or object in the external world. I use ‘proximal stimulus’ to refer to the property or object in the substituting modality (audition or touch), such as a particular tone delivered by a visual-to-audio SSD. By ‘distal stimulus’ I mean the property or object in the substituted modality (vision) that the proximal stimulus represents, such as the brightness represented by the tone.

(‘soundscapes’), which the user hears through headphones’ (Kiverstein, Farina, and Clark forthcoming). The vOICe represents the vertical spatial location of an object by using higher or lower pitched tones (the higher the tone, the higher the spatial location of the object). It represents horizontal spatial location by using the left and right headphone. If the spatial location of an object is to the left (or right), the vOICe plays a tone out of the left (or right) headphone. The vOICe represents the brightness of an object with volume (the brighter the object, the louder the volume).

Consider a subject using the vOICe. A subject using the vOICe has already been exposed to auditory stimuli, prior to her first experience with the SSD. Contrast this with an infant. For an infant, there is a time in which the sounds of birds chirping, voices talking, and fans whirring are all new to him. For the user of the vOICe, on the other hand, she has already heard all of these sounds and many others. Auditory stimuli are not new to her.

SSD users have been pre-exposed not just to auditory or tactile stimuli in general, but also to some of the specific proximal stimuli types delivered by these devices. Consider the EyeMusic, an SSD that is like the vOICe in converting visual images into sounds, but is unique in also delivering information about colour. The EyeMusic represents five colours with notes from five ordinary instruments like a piano (white) and a marimba (blue) (Levy-Tzedek et al. 2012b: 316). In other words, the device delivers proximal stimuli to which the user is undoubtedly familiar: the notes of a piano and a marimba. Due to latent inhibition, however, prior exposure interferes with new learning. So, new learning might be harder to accomplish with piano and marimba notes. To take a different example, consider ‘TVSS,’ which stands for ‘Tactile-vision sensory substitution.’ TVSS devices convert visual information picked up by a camera into tactile information, which is delivered to the skin (to the tongue or back, in

particular). As Bach-y-Rita and Kercel describe it, however, the information delivered by TVSS is ‘the usual cutaneous information (pressure, tickle, wetness, etc.)’ (2003: 543). But certainly pressure, tickles, and wetness are familiar stimuli, which have been experienced before in a different context. Once again, however, research on latent inhibition shows that prior exposure interferes with new learning. So, it might be harder to train someone with TVSS, given that the device delivers pressure, tickles, and wetness—all familiar stimuli.

It might be objected here that the context of, say, the notes of a marimba surrounded by other SSD stimuli is importantly different from the context of the notes of a marimba surrounded by the notes of other instruments in a band. The worry is that latent inhibition of the marimba notes in one context might not transfer into the SSD context. But latent inhibition does transfer in these contexts, as experiments on latent inhibition for complex stimuli show (for a good summary, see McLaren and Mackintosh 2000: 221). Latent inhibition affects the individual features of a complex stimulus, not on the complex stimulus as a whole. If a stimulus, like a marimba note, is presented with one group of stimuli in one context, but with another group of stimuli in another context, latent inhibition affects the marimba, and transfers across contexts.⁴

I have been arguing that prior exposure interferes with new learning. I now want to argue that there are still ways in which sensory substitution researchers can avoid prior exposure. They can start by determining which stimuli are novel to SSD users by using ERPs or by measuring pupil dilation. This might be a complicated process, since a token stimulus can fall under

⁴ A second objection based on context is that the environmental context in an SSD training situation is importantly different from the environmental context in which subjects have been previously exposed to piano notes, marimba notes, pressure, tickles, and wetness. The worry is that latent inhibition in the prior context might not transfer into the SSD context. There is indeed quite a bit of evidence for the context-sensitivity of the latent inhibition effect (for a good summary, see Hall 1991: 75-78). Roughly, a change in context often attenuates the latent inhibition effect. Crucially, however, context sensitivity is itself attenuated if the subject is already familiar with the prior context before she is exposed to the stimuli (McLaren et al. 1994; McLaren and Mackintosh 2000: 224). In other words, if an infant is familiar with her crib context, and is then exposed to pressure and tickles, these stimuli will be latently inhibited, and that effect can transfer into the SSD learning context.

multiple types. An mbira note, for instance, might be novel as an mbira note, but familiar as the note C-sharp. In such a case, a researcher would need to carefully determine which feature is novel in order to construct the most novel stimulus. After researchers have determined such stimuli, they can program SSDs such that the devices deliver them. Since the EyeMusic represents colours with notes from familiar instruments like a piano and a marimba, for instance, these could be switched to notes of more unfamiliar types of instruments (the notes of an mbira might be a good substitute for piano notes, for example). Perhaps one reason why a piano and a marimba were chosen is because they are instruments that are familiar to the researchers who selected them. But novelty, not familiarity, facilitates perceptual learning. To speed up the training, more obscure instruments might be a better choice. This is a testable hypothesis.

The EyeMusic also uses tones that make up a pentatonic scale (Levy-Tzedek et al. 2012a: 3). This makes sense, because it means that the tones will sound consonant together when heard by someone using the device. However, it is also at odds with efficient perceptual learning. After all, a chord composed from the pentatonic scale will likely be a familiar chord. A scale less familiar to most people, such as the Locrian scale, might deliver more efficient perceptual learning than the pentatonic scale. Of course, it is important to strike a balance between user friendliness and efficient perceptual learning. The pentatonic scale is certainly pleasurable to listen to, and that may be why it was chosen. To balance efficient perceptual learning with user friendliness, however, researchers might use a pleasurable sounding scale that is also unfamiliar to the majority of users. This issue, however, suggests a more general point: researchers will often need to strike a balance between several different desiderata (including efficient engineering), and some of them may be in conflict with efficient perceptual learning.

b. Harnessing Latent Inhibition for Efficient Generalization

Consider a second way that researchers can use latent inhibition to improve SSD learning. Users of SSDs become able to recognize new token objects of the same type on which they have been trained. So, for instance, if they were trained on a maple leaf with the vOICE, they can become able to recognize new instances of maple leaves using that SSD. This process is called ‘generalization’ (see Kim and Zatorre 2008, and Deroy and Auvray forthcoming).

Although peripheral to the main argument of the paper, one thing worth mentioning is that there will be more and less optimal ways to achieve generalization through latent inhibition.

Suppose that the task is to get an SSD user to recognize a novel maple leaf as a maple leaf when she sees it. First, consider the least optimal training scenario for this, given latent inhibition. Suppose that immediately after picking up the sensory substitution device for the first time, the user is exposed just to a maple leaf, but not to other kinds of objects. Given that scenario, each of the maple leaf’s features will be equally latently inhibited. This will make it harder to recognize a new token maple leaf as a maple leaf when presented with one. It will be harder to recognize this new maple leaf as such because its prototypical features are unlikely to stand out any more than any other feature. This is because all of its features are equally likely to be latently inhibited, both the prototypical and non-prototypical features alike. To enhance recognition, what we need is for the non-prototypical features to be latently inhibited, so that the prototypical features stand out in contrast.

Suppose, on the other hand, that a user is initially exposed an equal number of times to a wide range of objects different from maple leaves, but sharing some maple leaf features (they might be other kinds of leaves, or non-leaves with pointy edges, etc.). In that case, due to latent inhibition, the user will be less habituated to the features shared between maple leaves and non-maple leaves, resulting in those features becoming less likely to enter into associations, and so

less likely to be good candidates for use in learning. In contrast, the user will be less habituated to the features unique to the maple leaf (such as the unique pointy shape of its edges). When introduced to a new token maple leaf, its prototypical features are more likely to enter into associations. As such, those features are better candidates for use in learning.

2.2 Attentional Weighting

After latent inhibition, a second perceptual learning effect is ‘attentional weighting,’ the idea that the weight of attention can change by ‘increasing the attention paid to perceptual dimensions and features that are important, and/or by decreasing attention to irrelevant dimensions and features’ (Goldstone 1998: 588).⁵ When you learn to recognize a wren, for instance, you might increase attention to the shape of its body, while decreasing attention to the colour of the bird’s plumage. If so, you have changed your original attentional weighting. For experts, like a birdwatcher, part of expertise is developing the ability to attend to relevant perceptual dimensions of birds.

Which perceptual dimension the subject attends to is often a function of the relevance of the task (see Goldstone and Byrge forthcoming). For example, when picking the right book to slip into your bag when you are in a hurry, colour and size might be important dimensions to attend to. In contrast, weight might be more important than colour or size for picking out a box of books from other boxes when you are moving. It is important to recognize that attentional weighting has as much to do with what the subject chooses to ignore, as it does with what she chooses to attend to (Goldstone 1998: 588-89). When you attend to the weight of the box as you are moving, you are also choosing to ignore many other features of the box, even perhaps its colour and size.

⁵ My exposition of attentional weighting and unitization/differentiation is largely from Goldstone 1998 and Goldstone and Byrge forthcoming.

It might be tempting to think of attentional weighting as something that is easy to change. After all, you might just as easily attend to the colour or size of the box of books, rather than the weighting. However, there is also evidence that some kinds of attentional weighting become ingrained. As Goldstone points out, ‘attentional highlighting of information occurs even if it is to the detriment of the observer’ (1998: 589). For instance, letters used first as targets but later used as distractors to be ignored, still grab the attention of subjects (Goldstone cites Shiffrin and Schneider 1977). In such a case, an attentional weighting has become ingrained, and persists even if the subject wants to change it.

A second piece of evidence that attentional weighting is not always easy to change is the following. Through experience we develop dispositions to attend in a particular way. For instance, experienced soccer defenders focus longer on the opponent’s hips than less experienced players do (Williams and Davids 1998), while experienced goalkeepers fixate longer on the non-kicking leg than less experienced goalkeepers (Savelsbergh et al. 2002). To say that these attentional patterns are easy to change, however, would be to ignore the fact that these are habits developed through years of experience. Experienced soccer players are disposed to attend in a particular way, and that attentional pattern has become deeply ingrained, fortified by years of experience.

2.2.1 Attentional Weighting and Sensory Substitution

Unweighting and Reweighting Attention in Sensory Substitution

I want to offer a novel proposal for how researchers can weight attention in a way that facilitates SSD learning. First, however, let me say something about the effect of prior attentional habits on SSD learning. Contrast the learning situation for an SSD user with the learning that takes place in ordinary perceptual development. For an infant, the perceptual learning process

involves developing an attentional weighting, but does not require breaking any past attentional disposition. In contrast, an SSD user has already developed attentional habits. The learning that takes place in sensory substitution occur on stimuli to which the user is familiar, either as an auditory or tactile stimulus in general, or as a particular type of stimulus such as the notes of a piano or a marimba with the EyeMusic (Levy-Tzedek et al. 2012a; 2012b), or as ordinary tickles, pressure, and wetness to the skin with TVSS (Bach-y-Rita and Kercel 2003: 543). Training a subject on an SSD might be difficult since it involves breaking an attentional habit, as opposed to just weighting attention. Once again, novel stimuli become important for training on SSDs. If SSDs were to employ novel stimuli, then subjects might simply develop an attentional weighting for the SSD's perceptual array without having to first break their attentional habits.

Now let me suggest the novel way that researches can weight attention to facilitate SSD learning. In a 2013 study on perceptual learning, subjects were trained to classify different kinds of fish locomotion (Jarodzka et al.). Three groups were shown the same four videos, each with a single fish swimming a different locomotion pattern. Each video was accompanied by an audio instruction from a professor of marine zoology explaining the relevant features of each locomotion pattern (Jarodzka et al. 2013: 65). As the professor delivered his instructions, he was also watching each video, and his eye movements were tracked. A control group was shown the four videos with the accompanying audios. A second group was given the four videos with audio, but in addition, each video had a dot that tracked the spot where the professor's eye-movements were focusing. A third group was given the videos and audio, along with a spotlight which both highlighted the spot where the professor's eye movements were focusing and filtered out the other distracting details.

In the testing phase of the experiment, subjects were shown four new videos, one for each locomotion pattern, and they were rated for the efficiency at which they fixated on one of the identifying features of the locomotion pattern, as well as the length at which they fixated there. Those in the dot group fixated on the identifying features quicker and focused there longer than the control group. Those in the spotlight group were even quicker, and fixated even longer. Both the dot and the spotlight group were also better than the controls at interpreting what they saw when responding to multiple-choice questions after viewing videos.

The fish locomotion study suggests a novel way to enhance the learning process with sensory substitution devices: attentional cueing might serve a useful function during training. More specifically, cues could shift the weight of attention onto the more task-relevant stimuli delivered by the SSD, and away from less task-relevant stimuli. Suppose you are training subjects with the vOICe or EyeMusic, and the audio is correlated with a complicated pattern on a computer screen. An attention-grabbing auditory cue such as a horn might be placed just temporally prior to stimuli that researchers want to highlight. This would cause the viewer to focus on the subsequent stimulus. Just as the dot drew the attention of subjects in the fish locomotion study, the horn could draw the auditory attention of subjects when training on SSDs. Importantly, however, the results of the fish locomotion study suggest that an auditory *spotlight* would be even better. Suppose again that you are training subjects with the vOICe or EyeMusic, and the audio is correlated with a complicated pattern on a computer screen. White noise could be played over distracting details to dampen them, but not over important stimuli. Just as the spotlight highlighted important features of fish locomotion, the lack of white noise could highlight important features in a sensory substitution training session.

In short, very recent literature on perceptual learning suggests some novel techniques that researchers might use to teach someone how to use an SSD. These ‘sensory substitution training wheels,’ might give users a better training experience with the device, and make them more proficient with SSDs once the training wheels are removed.

2.3 Unitization and Differentiation

After latent inhibition and attentional weighting, a third perceptual learning effect is unitization/differentiation. In the process of ‘unitization,’ perceptual features that were once treated by the perceptual system as distinct are later treated by the system as a single unit. A whole host of objects have been shown to be first processed as distinct parts, and later processed as a unit. This is the case not just for cats and cars, but also for objects constructed in the lab. For instance, Gauthier and Tarr constructed a set of objects called ‘Greebles,’ which shared a set of spatial features in common. When the subjects were exposed to the Greebles for long enough, they would begin to process them as units (Gauthier and Tarr 1997).

In the process of ‘differentiation,’ perceptual features that were once treated by the perceptual system as unified are later treated by the system as distinct parts. For instance, William James discusses a case where a man becomes able to distinguish by taste between the upper and lower half of a bottle for a particular kind of wine (1890: 509). Assuming that the man’s perception itself has changed, not just his background belief about what his perception amounts to, James’ case is a case of perceptual differentiation. At first, that kind of wine was treated as gustatorily unified by the man’s perceptual system, while later the upper and lower half of a bottle are treated as gustatorily distinct.

Unitization and differentiation seem at first glance to be opposites. After all, while in unitization, features once treated as distinct are later treated as a single unit, in differentiation,

features once treated as a single unit are later treated as distinct (Goldstone 1998: 602). In fact, however, unitization and differentiation are really just performing the same task in two different ways (Goldstone and Byrge forthcoming). Specifically, both unitization and differentiation create new perceptual units. Goldstone puts the point in the following way: ‘[U]nitization and differentiation are both processes that build appropriate-sized representations for the tasks at hand’ (1998: 604). Both unitization and differentiation give us units to attend to, search for, and build up into more complex percepts (Goldstone and Byrge forthcoming). It’s just that in unitization, these units are composite wholes, while in differentiation the building blocks are parts.

2.1.1 Sensory Substitution and Unitization/Differentiation

Re-Unitizing and Re-Differentiating

As with latent inhibition and attentional weighting, by the time a subject first uses an SSD, her perceptual array will already be arranged in a particular way through the processes of unitization and differentiation. Some features will be arranged together as units, while other features will be differentiated from one another as parts. This means that for some features, unitization or differentiation will need to be undone.

As with latent inhibition and attentional weighting, one way to avert much of the unlearning process for unitization and differentiation is to program SSDs such that they deliver novel stimuli to users. While familiar stimuli may have undergone differentiation and unitization, novel stimuli will not have undergone either of these perceptual processes. As a result, the use of novel stimuli might circumvent much of the unlearning process, making the perceptual learning process more efficient as a whole. After all, novel stimuli will require only unitization and differentiation, without requiring first that these processes are undone from the prior context.

3. A Philosophical Conclusion

I want to close with a philosophical upshot. Very recently, the discussion on SSDs has revolved around the nature of one's perceptual experience once one has fully integrated an SSD. In particular, the question is the following: Should that perceptual experience be classified in the substituting modality (as auditory or tactile), in the substituted modality (as visual), or in a new modality? Enactivists such as Noë, Hurley, and O'Regan have argued that the perceptual experience should be classified as visual (Hurley and Noë 2003; Noë 2004; and O'Regan 2011). Block has argued that it should be classified as auditory or tactile (2003). Kiverstein, Farina, and Clark have argued that the perceptual experience should be classified as part of a new modality (forthcoming). Deroy and Auvray have argued that these three options are insufficient (forthcoming). Assuming (pace Deroy and Auvray) that the three options are sufficient, my argument suggests a novel way that we can answer this question empirically.

If perceptual experience with an SSD is in the substituted modality (vision), or in a new hybrid modality, then think about the stimuli involved in that experience. Assuming that the SSD user is blind from birth, if the perceptual experience is visual, then the stimuli will be novel to the subject. Likewise, if the perceptual experience is in a new hybrid modality, then the stimuli will also be novel to the subject. But remember that there are ways to test whether a stimulus is novel or familiar to a subject. One option is to examine brain responses. Novel stimuli activate networks in the brain that have a signature and are detectable. This is done by using electrophysiological recordings of scalp and intracranial event-related potentials (ERPs) (see Knight, 1996). Another option, for cortically blind subjects at least, is to examine their pupil response to determine novel stimuli, since the pupil dilates more for novel stimuli (Naber et al. 2013), and there is evidence that this occurs even for non-visual perceptual stimuli (Naber et al.

2013: 2). By using ERPs or by measuring pupil dilation, researchers might determine whether sensory substitution stimuli are novel to a subject. If they are novel, then we have evidence that the perceptual experience with an SSD is best classified either in the substituting modality (as visual), or in a new hybrid modality. If the stimuli are familiar, then we have evidence that the perceptual experience remains in the substituted modality (as auditory or tactile).

There are a lot of remaining questions about perceptual learning and SSDs. First, the learning involved with SSDs has a motor component. Subjects need to learn how to control the camera with the vOICE, for instance. The motor component of sensory substitution learning is an important topic, which I have not addressed here. Second, although this paper was about sensory substitution and perceptual learning, it focused predominantly on what perceptual learning can tell us about sensory substitution. A lot can still be said about what sensory substitution can tell us about perceptual learning. Finally, in discussing what perceptual learning can tell us about sensory substitution, the paper focused largely on the proximal stimuli delivered by SSDs, and while it focused a little bit on the distal stimuli, that was not the main focus of the paper. A lot can still be said about what research on perceptual learning can tell us about the distal stimuli that SSDs track.

References:

- Bach-y-Rita, P., & W Kercel, S. (2003), 'Sensory substitution and the human-machine interface', *Trends in cognitive sciences*, 7(12), 541-546.
- Block, N. (2003), 'Tactile sensation via spatial perception', *Trends in Cognitive Sciences*, 7, 285-286.
- Deroy, O., & Auvray, M. (forthcoming), 'Beyond vision: The vertical integration of sensory substitution devices', in Stokes, Biggs, & Matthen (eds.) *Perception and its Modalities*, Oxford University Press.
- Gauthier, I., & Tarr, M.J. (1997), 'Becoming a "greeble" expert: Exploring mechanisms for face recognition', *Vision Res.* 37: 1673-1682.
- Goldstone, R. L. (1998), 'Perceptual learning', *Annu. Rev. Psychol.* 49, 585-612.
- Goldstone, R. L. & Byrge, Lisa A. (forthcoming), 'Perceptual Learning', in Mohan Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*, Oxford University Press.
- Hall, G. (1991), *Perceptual and associative learning*, (Oxford, Oxford University Press, Clarendon Press).
- Held, R., Ostrovsky, Y., de Gelder, B., Gandhi, T., Ganesh, S., Mathur, U., & Sinha, P. (2011), 'The newly sighted fail to match seen with felt', *Nature neuroscience*, 14(5), 551-553.
- Hurley, S., and Noë, A. (2003), 'Neural plasticity and consciousness', *Biology and Philosophy*, 18, 131-168.
- James, W. (1890), *Principles of Psychology*, (New York, Holt).
- Jarodzka, H., van Gog, T., Dorr, M., Scheiter, K., & Gerjets, P. (2013), 'Learning to see: Guiding students' attention via a Model's eye movements fosters learning', *Learning and Instruction*, 25, 62-70.
- Kim, J-K, & Zatorre, R. J. (2008), 'Generalized learning of visual-to-auditory substitution in sighted individuals', *Brain Research*, 1242, 263-275.
- Kiverstein, J., Farina, M., and Clark, A. (forthcoming), 'Substituting the Senses', in Mohan Matthen (ed.), *The Oxford Handbook of the Philosophy of Perception*, Oxford University Press.
- Knight, Robert T (1996), 'Contribution of human hippocampal region to novelty detection', *Nature* 383, 6597: 256-259.

- Kupers, R., & Ptito, M. (2004, August), ‘‘Seeing’ through the tongue: cross-modal plasticity in the congenitally blind’, in *International Congress Series*, Vol. 1270, 79-84, Elsevier.
- Levy-Tzedek, S., Novick, I., Arbel, R., Abboud, S., Maidenbaum, S., Vaadia, E., & Amedi, A. (2012a), ‘Cross-sensory transfer of sensory-motor information: visuomotor learning affects performance on an audiomotor task, using sensory-substitution’, *Scientific reports*, 2.
- Levy-Tzedek, S., Hanassy, S., Abboud, S., Maidenbaum, S., Amedi, A. (2012b), ‘Fast, accurate reaching movements with a visual-to-auditory sensory substitution device’, *Restorative Neurology and Neuroscience* 30, 313-323.
- Lubow, R.E. (2010), ‘A short history of latent inhibition research’, in *Latent Inhibition Cognition, Neuroscience and Applications to Schizophrenia*, Robert Lubow & Ina Weiner (eds.), Cambridge University Press.
- Lubow, R.E., & Moore, A.U. (1959), ‘Latent inhibition: The effect of non-reinforced preexposure to the conditional stimulus’, *Journal of Comparative and Physiological Psychology*, 52, 415-419.
- McLaren, I.P.L., Kaye, H. and Mackintosh, N.J. (1989), ‘An associative theory of the representation of stimuli: applications to perceptual learning and latent inhibition’, in R.G.M. Morris (ed.) *Parallel Distributed Processing - Implications for Psychology and Neurobiology*, (Oxford: Oxford University Press).
- McLaren, I. P. L., Bennett, C., Plaisted, K., Aitken, M., and Mackintosh, N. J. (1994), ‘Latent inhibition, context specificity, and context familiarity’, *The Quarterly Journal of Experimental Psychology*, 47(4), 387-400.
- McLaren, I.P.L. and Mackintosh, N.J. (2000), ‘An elemental model of associative learning: I. Latent inhibition and perceptual learning’, *Animal Learning and Behavior*, 38(3), 211-246.
- Naber, M., Frassle, S., Rutishauser, U., & Einhauser, W. (2013), ‘Pupil size signals novelty and predicts later retrieval success for declarative memories of natural scenes’, *Journal of Vision*, 13(2):11, 1–20.
- Noë, A. (2004), *Action in perception*, (Cambridge, MA: MIT Press).
- O’Regan, J. K. (2011), *Why red doesn't sound like a bell: Understanding the feel of consciousness*, (Oxford: Oxford University Press).
- Ptito, M., Moesgaard, S. M., Gjedde, A., & Kupers, R. (2005), ‘Cross-modal plasticity revealed by electrotactile stimulation of the tongue in the congenitally blind’, *Brain*, 128(3), 606-614.

Savelsbergh, G.J.P., Williams, A.M., Van der Kamp, J., & Ward, P. (2002), 'Visual Search, Anticipation and Expertise in Soccer Goalkeepers', *Journal of Sports Sciences*, 20, 279-287.

Shiffrin RM, Schneider W. (1977), 'Controlled and automatic human information processing: II. Perceptual Learning, automatic attending and a general theory', *Psychol. Rev.* 84:127-90.

Tolman, E. C., & Honzik, H. C. (1930), 'Introduction and removal of reward, and maze performance of rats', *University of California Publications in Psychology*, 4, 257-275.

Williams, A. M. and Davids, K. (1998), 'Visual Search Strategy, Selective Attention, and Expertise in Soccer', *Research Quarterly for Exercise and Sport*, 69, 111-128.