

## Sensory Substitution and Non-Sensory Feelings

David Suarez, Diana Acosta Navas, Umut Baysan, and Kevin Connolly

**Abstract:** One of the central limitations of sensory substitution devices (SSDs) is their inability to reproduce the non-sensory feelings that are normally associated with visual experiences, especially hedonic and aesthetic responses. This limitation is sometimes reported to cause SSD users frustration. To make matters worse, it is unclear that improvements in acuity, bandwidth, or training will resolve the issue. Yet, if SSDs are to actually reproduce visual experience in its fullness, it seems that the reproduction of non-sensory feelings will be of some importance. We offer a novel solution. Researchers can produce hedonic and aesthetic responses by eliciting these feelings artificially, pairing distal objects that should be pleasurable to pleasurable stimulus outputs from the SSD. We outline two strategies for accomplishing this: first, by means of a prefixed, hardwired, association of pleasant distal objects to pleasant stimulus outputs from the SSD; second, by means of a flexible, feedback-based association which creates associations based on a subject-directed matching of distal objects to patterns of stimuli from the SSD which the subject takes to have the corresponding hedonic properties. We evaluate some problems with both strategies, and we argue that the feedback-based strategy is more promising. Researchers can use this strategy to help the blind, allowing them to take pleasure in the objects they perceive using SSDs.

### 1. Introduction

Perception provides us with sensory contents like shapes, colours, flavours, and smells. In ordinary perception, these contents are accompanied by a variety of non-sensory feelings. Perception arouses *noetic feelings* concerning the presence or familiarity of what one is perceiving, *emotions* such as love and hate, *hedonic responses* such as pleasure and pain, and *aesthetic feelings* of beauty and ugliness. Although some non-sensory feelings seem to be naturally associated with specific sensory contents, non-sensory feelings do not seem to be supervenient on the sensory contents they accompany, since non-sensory feelings can vary while sensory contents remain the same. In Capgras syndrome, for instance, patients perceive the face of a close relative, yet do not have a feeling of familiarity. Such patients consider their relatives to be impostors, since they have no feeling of familiarity when seeing them (see Christodoulou, 1977).

The failure of supervenience of non-sensory feelings on sensory contents seems also to be supported by the cases in which sensory substitution devices (SSDs)<sup>1</sup> are used. SSD users very often report an absence of hedonic, emotional and aesthetic feelings, despite being capable of identifying the visual scenes that would typically elicit those feelings. Hedonic responses are generated by some stimuli in the substituting modality<sup>2</sup>, but these responses are not associated to the visual scenes that would normally generate them. Therefore, even though many users of SSDs are eventually able to attribute stimuli in substituting modalities to distal objects, it takes training—and, in many cases, explicit instruction in the mapping between substituting and substituted modalities—to generate the right non-sensory feelings from the appropriate distal objects. Ordinary distal attribution occurs “when the sensations felt on the skin or in the ears are projected onto a distant perceived object, which should correspond to the one captured by the camera” (Deroy and Auvray, forthcoming). After a period of training, subjects come to attribute their sensory contents to distal objects rather than to the proximal stimuli delivered by the SSD; but this is not (usually) the case with non-sensory feelings.

This raises the question whether an analogous phenomenon to distal attribution can occur with *non-sensory feelings*, given enough and appropriate training. So, for example, will subjects project their hedonic responses onto the distal objects perceived by means of the SSD, rather than attributing them to the outputs of the SSD itself? The degree to which this will occur is an empirical question, to be answered by future research, but here we should note two things. First,

---

<sup>1</sup> A sensory substitution device (SSD) delivers information about the environment, normally perceived through stimulation in one sensory modality (the “substituted modality”), through the production of stimulation in another sensory modality (the “substituting modality”). Most SSDs aim to substitute for vision, and consist of a video camera that feeds information into a conversion unit which then converts that information into auditory or tactile stimuli. Such devices are often used by the blind to assist in their autonomous navigation of the world.

<sup>2</sup> We take a stimulus to be an objective property or object in the external world. We 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” we mean the property or object in the substituted modality (vision) that the proximal stimulus represents, such as the brightness represented by the tone.

distal attribution of non-sensory contents is part of normal experience. When we see a beautiful painting, we project the pleasure we feel onto the painting; it is not just our immediate visual sensations that we take to pleasing, but the painting itself, which is pleasurable to see. Second, for this very reason, distal attribution of non-sensory contents is both expected and desired by SSD users.

One might wonder why SSDs *should* duplicate normal non-sensory feelings. After all, if sensory substitution aims only at providing alternative means for delivering sensory contents, then the duplication of *non*-sensory feelings would be irrelevant to that goal. However, the duplication of non-sensory feelings can plausibly be included amongst the *therapeutic* goals of sensory substitution. Beyond enabling practical capacities dependent on sensory contents such as object recognition and autonomous navigation, sensory substitution also has the potential to provide patients with pleasure, emotional connection, aesthetic experience, and a sense of being grounded in a shared reality. As Bach-y-Rita and his colleagues (2003; 2002) report, users of TVSS devices<sup>3</sup> found the “absence of qualia” (i.e., the absence of affective responses) rather disappointing, if not disturbing. This suggests that users of SSDs hope for more than just enhancements to their practical capacities, but also for experiences that enrich their lives in other ways. Similarly, Spence (2013) has argued that flavor and olfactory substitution devices would be far less appealing if they failed to produce the pleasures associated with tasty or fragrant stimuli. In the absence of such pleasures, it seems unlikely that patients with impairments of taste or olfaction would have much of a desire for hedonically-neutral, purely informational substitutes. So, while sensory substitution does aim at enabling practical capacities, it seems

---

<sup>3</sup> Tactile-vision sensory substitution (TVSS) devices are SSDs that convert patterns of luminance picked up by a camera into isomorphically-organized tactile stimuli which are delivered to the skin through a matrix of solenoids usually mounted on the back, or a matrix of electrodes held on the tongue.

plausible to think that the provision of other life-enriching experiences constitutes a worthwhile therapeutic goal.

The goal of this paper is to explore whether or not sensory substitution devices can capture the non-sensory feelings that typically accompany perception. In section two, we present evidence that SSDs do not currently capture such feelings. In section three, we explore whether this failure is fully explicable by the fact that the acuity and bandwidth of the SSDs are poor. In section four, we offer a suggestion: by using a feedback mechanism, subjects might pair pleasurable objects in their environment to pleasurable sounds (or pleasurable patterns of touch) delivered by SSDs. This will enable them to experience hedonic responses when they perceive those objects through the SSDs.

## **2. Current SSDs fail to generate normal non-sensory feelings**

Since the early days of research on sensory substitution, it has been reported that SSDs do not always preserve the non-sensory feelings that are associated with normal perceptual experiences. In particular, SSDs fail to generate normal emotional and hedonic responses (Bach-y-Rita, 2002) and convey to the subject the “values and the quality of lived experience associated with perceived entities” (Lenay et al., 2003). The researchers call these absent features *qualia*. Evidence about the absence of (normal) qualia and the subsequent frustration in SSD usage include a case where two (presumably heterosexual) male subjects using a TVSS apparatus were shown pictures of nude women, and a case where subjects were shown pictures of their partners. In these cases, as Bach-y-Rita reports, although subjects were able to recognize what they were being shown, they were surprised and disappointed to find that they did not find these stimuli pleasing (2002, p. 509). Bach-y-Rita also reports a few cases in which ‘qualia’ were developed, but he emphasizes their scarcity (p. 510). Concerning the reports about the absence of (normal)

qualia in SSD usage, Lenay et al (2003) suggest that the subsequent frustration could result from the fact that blind subjects typically use sensory substitution devices alone rather than in groups. Lenay et al. argue that qualia are “closely linked to the existence of a shared history and collective memory, a memory which can only emerge in the course of interactions between several subjects in a common environment” (p. 285).

The disappointment experienced by SSD users is akin to the frustration experienced by people with early blindness who regain their vision after sight recovery operations. It should be noted that such cases are very rare (Ostrovsky et al., 2006), so it is difficult to draw strong conclusions from the available reports. Moreover, only a small number of these reports treat issues related to the experience of qualia. Nevertheless, one exception is the case of S.B. (Gregory & Wallace, 1963; Gregory, 1990). S.B. lost his sight when he was 10 months old, and regained it through a surgical intervention when he was 52 years old. Although he learned to recognize objects very successfully thereafter (he was even able to read the time from a clock a few days after the surgery), his perceptual experiences remained qualitatively different from those of a normally-sighted subject. Facial expressions that he was able to recognize did not carry emotional messages; he did not dislike any colors; he was not very enthusiastic about vivid and colorful scenes (he preferred to sit in the dark); and he was bored on sightseeing tours. While these examples do suggest that there were some non-sensory feelings associated with his visual experiences, his experiences were, nevertheless, clearly quite different from those of the normally-sighted. In particular, the difference between S.B.’s postoperative visual experience and a normally-sighted person’s visual experience could be interpreted as primarily a difference in his profile of hedonic responses. The difficulty of generating normal profiles of hedonic responses after sight recovery operations resembles the difficulty of transferring hedonic aspects

from one sensory modality to another. With respect to SSDs in particular, Deroy and Auvray note that “shapes perceived in one sensory modality are not directly associated to pleasures or pains felt while perceiving the same shape in another sensory modality” (2012, p. 4).

Bach-y-Rita (2002), however, is rather optimistic about the possibility of conferring such hedonic messages to the substituting modalities. He thinks that the case is analogous to the acquisition of a second language: “The emotional aspects of the new language are often lacking, especially with emotionally charged words and expressions, such as curse words. It appears that both spoken language and other sensory messages require long experience within the context of other aspects of cultural and emotional development to be able to contain qualia” (Bach-y-Rita, 2002, p. 510). So, with sufficient training, Bach-y-Rita suggests that SSD usage would generate normal hedonic responses. He also speculates that if such SSD “systems are used from infancy, the qualia will develop, even if the system is not continuously in use” (Bach-y-Rita, 2002, p. 510). In fact, data seem to support this hypothesis as “[s]ystems for blind babies ... have already provided some suggestive evidence for the development of qualia, such as the infant’s smile upon perceiving the mother’s approach” (Bach-y-Rita, 2002, p. 510).

Regardless of whether Bach-y-Rita’s optimism about capturing the normal qualitative aspects of experience through longer exposure to SSDs is justified or not, it can be fairly stated that, currently, SSDs fail to generate normal non-sensory feelings. In the following sections, we will consider the possible sources of the problem and discuss solutions to it.

### **3. Is the Problem Just Bandwidth or Acuity?**

In a 2001 study, Sampaio, Maris, and Bach-y-Rita took subjects who had received some training with TVSS, and tested them with a standard ophthalmological test. The “visual” acuity of these subjects measured on average only 20/430. Without training, their acuity with TVSS

was even worse, averaging only 20/860 (Sampaio, Maris & Bach-y-Rita, 2001, p. 205). Given these results, perhaps the reason why SSDs do not generate affective characteristics is because the acuity of a subject is so poor when using a SSD. That would explain why Bach-y-Rita's male SSD subjects did not respond in the expected way when presented with the picture of the nude women. The "picture" was not clear enough.

Vision and audition may present a bandwidth problem for SSDs, but it is not clear that the same would be true with respect to touch, smell, or taste. According to Charles Spence, we can transduce 40 bits per second in vision and 30 through hearing. In contrast, we transduce just 5 bits per second through the skin, and just 1 bit per second in taste and smell (Spence, 2013). If the bandwidth of vision and audition is so high that it cannot be fully captured by SSDs making use of the remaining modalities, then this would explain why users of visual or auditory SSDs would have poor acuity while using them, and hence lack the expected non-sensory responses. By contrast, the bandwidths of touch, smell, and taste do not present a comparable problem. This is because the bandwidths of these senses are low enough to be captured by an SSD using vision or audition as the substituting modality. But take a sense like smell. Even though smell does not present the bandwidth problem that vision presents, it seems plausible that an SSD designed to substitute for olfaction would still not capture the non-sensory feelings associated with smell. For example, suppose we were able to deliver all of the information relevant to olfactory perception through vision instead, with the same informational bandwidth, and the same acuity as olfaction. Nevertheless, it seems plausible to think that delivering this information would still be insufficient, on its own, to generate the pleasure many of us associate with the smell of a rose. Given this plausible assumption, it can be suggested that receiving the information normally provided by a given sensory modality may not be sufficient (even with the appropriate SSD

acuity, bandwidth, and training) for the duplication of the substituted modality's normal profile of non-sensory feelings.

#### **4. Natural vs. Artificial Associations**

Even if SSDs are currently incapable of reproducing accurate profiles of non-sensory feelings, it is worth considering whether it would be possible to create artificial associations between the distal stimuli that SSDs are tracking, and the natural hedonic responses SSDs produce with proximal stimuli. For example, users of the FeelSpace belt<sup>4</sup> report having pleasurable sensations when using the device (Konig, 2013), and users of the vOICe<sup>5</sup> describe some of the sounds they hear as beautiful. Researchers might take these natural responses and artificially associate them with pleasant or unpleasant distal visual scenes. The problem is that these natural responses are not associated to the distal stimuli that SSDs are tracking, but rather to the proximal auditory or tactile stimuli produced by the SSDs. As a result, when using SSDs, subjects' natural responses to perceiving objects rarely correspond to their expectations about what those responses should be. When subjects "see" objects that they expect to elicit hedonic responses, but do not experience what they expect, they often feel frustrated. One question, then, is whether it is possible to create artificial associations between a) objects that normally elicit hedonic responses in sighted subjects, and b) proximal stimuli that SSD users find naturally pleasurable. For example, can researchers produce SSDs that would associate a pleasant looking lilac with a naturally pleasant sound or pattern of touch? Would this enable researchers to reproduce normal profiles of non-sensory feelings?

---

<sup>4</sup> The FeelSpace belt is a belt designed to help subjects orient in space by indicating the magnetic North through vibrations around subjects' waist.

<sup>5</sup> As Kiverstein, Farina, and Clark describe it, "The vOICe is a visual-auditory substitution device that works by transforming images from a digital camera embedded in a pair of sunglasses into auditory frequencies ("soundscapes"), which the user hears through headphones" (forthcoming).



Furthermore, supposing all of this were successful, another issue remains unsettled: whether subjects will attribute these non-sensory feelings to the objects perceived. So, for example, assume that an SSD is successful in associating the presentation of a pleasant-looking lilac with a naturally pleasant sound. If we are aiming to duplicate *all* the aspects of a normal visual experience then something more than a congruence of a pleasant object with a pleasant proximal stimulus may be required—namely, that the subject experiences the pleasure associated with their perception as deriving *from the lilac itself*, rather than from the sound produced by the SSD. In what follows, we hope to show a way in which it is possible to create artificial associations, but we leave it as an open empirical question whether distal attribution of non-sensory feelings will occur.

With regard to whether it is possible to create artificial associations, there seems to be a relevant difference between devices that deliver visual information by means of sound and those that use touch. Unlike auditory stimuli, tactile stimuli can only vary along two dimensions: intensity and spatiality. This makes it harder to devise patterns of touch that are more pleasurable than others, in a way that resembles the wide spectrum of hedonic responses that visual images are capable of eliciting. On the other hand, it is known that some sound patterns are more pleasant than others. Both musical theory and psychological research about emotional responses to sounds could be consulted so as to find typically pleasing sounds. Assuming that such reliable information about the hedonic qualities of sounds is available, would it also be possible to use the sounds that we know to be pleasant or emotionally evocative to convey pleasant or emotionally evocative visual stimuli to blind subjects?

There are two possible ways to do this: either by means of a prefixed, hardwired, association of pleasant distal objects to pleasant proximal stimuli; or by means of a flexible,

feedback-based mechanism, which creates associations based on subjects' matching of the hedonic properties of distal objects to patterns of proximal stimuli that SSD users take to have the corresponding hedonic properties.

The first option would involve using the current code used by the vOICE to "translate" images into sounds, and using the available knowledge of natural cross-modal correspondences to complement this code. This would involve figuring out which sounds are naturally associated to the distal objects that produce certain hedonic responses, and adding those sounds to the sound patterns that represent those distal objects. Another, perhaps simpler, way to do this would be to select sounds that typically cause hedonic responses (regardless of the objects to which they're typically associated), and add these sounds to the system. This would involve programming the device in such a way that it is capable of discriminating certain visual patterns as being, for example, either a face or a landscape. With this ability for categorization, the device could associate types of pleasant sounds to categories of visual scenes, and play those sounds, in addition to the original code, when the subject is confronted with one of those scenes. This association would be programmed onto the device, just like the basic code.

Hardwiring pleasant acoustic patterns into an SSD is a way to secure that seeing a pleasant object will elicit the expected hedonic responses. This option, however, seems artificial, distorting the normal phenomenology of pleasure. That is, this strategy would impose fixed hedonic responses onto the subject, while failing to account for the fact that, in the normal case, hedonic responses are flexible and dynamic. For instance, many hedonic responses are acquired only after several exposures to, say, a painting or a person's face. Also, the same visual scene doesn't always elicit the same hedonic response. This variation is dependent on a wide variety of factors. Some, like lighting conditions, which a device could detect, could in principle be taken

into account in the programming. Others, like mood changes, are more arbitrary and highly subjective, and would be very difficult, if not impossible, to build into the device's architecture.<sup>6</sup>

Perhaps a more appropriate way of building these hedonic features into the SSDs is through a more flexible mapping between pleasant sounds and pleasant images, which is responsive to the subject's preferences. This task would require a feedback-based mechanism, through which the subject is able to set and alter associations between images and sounds. Consider how such a mechanism might have worked when Bach-y-Rita's male subjects using TVSS were shown pictures of nude women. First, the subject "sees" the woman. Then he notes the absence of a hedonic response. Next, he uses the feedback mechanism to indicate that the image should be pleasurable. The next time he "sees" that image, the device will use that feedback to deliver pleasurable proximal stimuli. The mechanism could include a "scale of pleasantness" for subjects to qualify visual scenes; or even allow the subject to classify visual scenes according to which kind of hedonic response they expect them to elicit.

The possibility of building such a mechanism raises two related issues. In the first place, building in a feedback mechanism could imply an elimination of the fixed mapping of visual-to-auditory properties, which is the very basis of the system, and the feature in virtue of which subjects can learn to use SSDs to recognize, discriminate, and perform other perceptual tasks. It is important to note that building in a feedback mechanism would change the nature of training with the device in a crucial way. Specifically, training would run in two directions, since both the subject and the device would learn from one another: subjects would learn the translation code of the device, in order to discriminate sensory contents; and through their feedback, devices would also "learn" to associate pleasant acoustic stimuli to what the subject considers to be pleasant

---

<sup>6</sup> This becomes even more complicated if we consider not only the variations in hedonic responses to images, but also the variations in hedonic responses to sounds. Since these two responses need not vary coherently, the interactions between them would become hard to deal with when programming such a device.

scenes. Whether this flexibility would hinder the learning process, or perhaps facilitate it, is an empirical question. The success of this mechanism, however, would depend on the possibility of maintaining a fixed code for the translation of visual information into auditory information, while having some flexibility for pleasant sounds to be assigned to the visual scenes that the subject considers pleasant.

On the other hand, introducing this flexibility may be a good strategy to address the problems of artificiality and rigidity that would appear if we had hardwired associations between pleasant sounds and pleasant sights. Yet, it is not entirely clear whether it would be a good strategy to reproduce the way in which visual images produce hedonic responses. Having such a feedback allows for the possibility of having hedonic responses that aren't imposed on the subject, but it might cause the opposite problem: the hedonic response would appear more like a choice on the part of the subject, rather than a response to a pleasant image.

Nevertheless, it could be the case that, after a period of training, the visual images themselves start to be perceived as pleasant. As mentioned above, Bach-y-Rita thought that hedonic qualities may appear in sensory substitution in the same way that they appear for the non-native speakers of a language: after a long exposure to the context and cultural aspects of the language (Bach-y-Rita, 2002, p. 510). This analogy also suggests that it is the use of the device which creates associations of visual scenes with hedonic responses, just as it is the use of the language that furthers the creation of such associations.

It may be thought that, when people learn to use a second language, the emotional aspects of the language are learned by two different means: first, by perceiving the way in which native speakers use emotionally charged words in particular situations; and second, by using the words in the appropriate contexts. If this analogy is correct, then it would seem that the development of

hedonic and emotional responses for users of SSDs would not only depend on the amount of training, but also on their active involvement in the training process. Non-native speakers learn to associate emotional contents to words, not just by hearing the words, but also by actively using them in emotionally charged situations. Initially, the emotional content may seem dissociated from the word; it may seem to the subject as if he is choosing to pair a neutral word with an emotionally charged situation. Nonetheless, after a period of training, the emotional content may be associated to the word itself, and not merely to the context in which it is used. Likewise, providing feedback to the SSD may constitute an active involvement of the user, in virtue of which he is able to couple hedonic responses with visual scenes that initially seem hedonically neutral. This active involvement might have a similar effect: after a while, the visual scenes themselves may come to seem pleasant. At some point, it may happen that subjects attribute the hedonic properties to the distal visual scenes, and not to proximal auditory stimuli.

This discussion opens a set of related philosophical questions: Is it possible for SSD users to associate pleasant auditory stimuli to pleasant visual scenes by giving feedback to the devices? Is it possible that, after a period of training, such an association leads to the attribution of pleasant properties to a distal visual scene, and not merely to the proximal auditory stimuli, or would the pleasantness remain attributed to the subject's choice instead of being attributed to the scene? Would the active involvement of the subject in the process of association make a difference to the way that pleasantness is experienced? More generally, would a feedback mechanism facilitate (some aspects of) training with SSDs, or would it hinder the learning process? These questions, and others in their neighborhood are open both to philosophical debate and to empirical research.

The main advantage of the feedback mechanism is that it could contribute to making the experience of SSD users more pleasurable. If the frustration reported by SSD users is not due to the absence of hedonic responses to the stimuli given to them (since there are hedonic responses to some sounds), but rather to the lack of consistency between their expectations of what should cause a hedonic response (e.g. seeing their wife's face) and what actually causes such a response (e.g., a meaningless pattern of sounds), this proposal should be helpful in addressing that problem. By building a feedback mechanism into the system, subjects would be given the possibility of pairing pleasant auditory stimuli to visual stimuli that they expect to be pleasant, thereby increasing the consistency of their expectations with their actual experiences. This way, their experiences would be more fluent, and this would, in turn, increase the likelihood of generating the expected non-sensory feelings by means of SSDs. We hope that the implementation of feedback mechanisms in SSDs will turn out to be an avenue of research that is worth pursuing.<sup>7</sup>

---

<sup>7</sup> We thank Susanna Siegel for her helpful comments on an earlier draft of the paper.

## References:

- Bach-y-Rita, P. (2002). Sensory substitution and qualia. In A. Noë & E. Thompson (Eds.), *Vision and mind: selected readings in the philosophy of perception* (pp. 497–514). Cambridge, MA: MIT Press.
- Bach-y-Rita, P., Tyler, M. E., & Kaczmarek, K. A. (2003). Seeing with the brain. *International Journal of Human-Computer Interaction*, 15(2), 285–295.
- Christodoulou, G. N (1977). The Syndrome of Capgras. *British Journal of Psychiatry*, 130: 556–564.
- Deroy, O., & Auvray, M. (2012). Reading the world through the skin and ears: a new perspective on sensory substitution. *Frontiers in Psychology*, 3(457), 1–13. doi:10.3389/fpsyg.2012.00457
- Deroy, O., & Auvray, M. (forthcoming). “Beyond vision: The vertical integration of sensory substitution devices.” In Stokes, Biggs, and Matthen (eds.) *Perception and its Modalities*, Oxford University Press.
- Gregory, R. L. & Wallace, J.G. (1963). Recovery from early blindness: A case study. *Quarterly Journal of Experimental Psychology*, Monograph 2. Heffers: Cambridge.
- Gregory, R. L. (1990). *Eye and brain: The psychology of seeing*. Oxford, UK: Oxford University Press.
- 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.
- Konig, Peter. (2013, March 26). *Properties and Mechanisms of Sensory Augmentation*. Presented at the Sensory Substitution and Augmentation Conference, The British Academy, London, UK.
- Lenay, C., Gapenne, O., Hannequin, S., Marque, C., & Genouelle, C. (2003). Sensory substitution: limits and perspectives. In Y. Hatwell, A. Streri, & E. Gentaz (Eds.), *Touching for Knowing* (pp. 275–292). Amsterdam: John Benjamins.
- Ostrovsky, Y., Andalman, A., Sinha, P., (2006). Vision following extended congenital blindness. *Psychological Science*, 17, 1009-1014.
- Sampaio E., Maris S., Bach-y-Rita P. (2001). Brain Plasticity: “Visual” Acuity of Blind Persons Via the Tongue. *Brain Research*, 908, 204-207.

Spence, C. (2013, March 26). *Commentary on Jonathan Cohen's "On Some Limitations of Sensory Substitution."* Presented at the Sensory Substitution and Augmentation Conference, The British Academy, London, UK.