SOUND-COLOUR SYNAESTHESIA: TO WHAT EXTENT DOES IT USE CROSS-MODAL MECHANISMS COMMON TO US ALL?

Jamie Ward, Brett Huckstep and Elias Tsakanikos

(Department of Psychology, University College London, London, UK)

Abstract

This study examines a group of synaesthetes who report colour sensations in response to music and other sounds. Experiment 1 shows that synaesthetes choose more precise colours and are more internally consistent in their choice of colours given a set of sounds of varying pitch, timbre and composition (single notes or dyads) relative to a group of controls. In spite of this difference, both controls and synaesthetes appear to use the same heuristics for matching between auditory and visual domains (e.g., pitch to lightness). We take this as evidence that synaesthesia may recruit some of the mechanisms used in normal cross-modal perception. Experiment 2 establishes that synaesthetic colours are automatically elicited insofar as they give rise to cross-modal Stroop interference. Experiment 3 uses a variant of the cross-modal Posner paradigm in which detection of a lateralised target is enhanced when combined with a non-informative but synaesthetically congruent sound-colour pairing. This suggests that synaesthesia uses the same (or an analogous) mechanism of exogenous cross-modal orienting as normal perception. Overall, the results support the conclusion that this form of synaesthesia recruits some of the same mechanisms used in normal cross-modal perception rather than using direct, privileged pathways between unimodal auditory and unimodal visual areas that are absent in most other adults.

Key words: synaesthesia, music, colour, cross-modal, audiovisual, ventriloquism

INTRODUCTION

For some individuals, auditory sensations automatically elicit a definite and reliable visual experience. For example, middle 'C' on a piano might be red but the note three octaves higher might be green (Ginsberg, 1923) and a "sharp rap on a table [may arouse] a faint but distinct narrow flash of brightness" (Whipple, 1900, p. 393). These individuals probably have sound-colour synaesthesia (or music-colour synaesthesia, also-known-as chromaesthesia; Cytowic, 2002). This type of synaesthesia is the focus of the present investigation.

Three experiments are reported which establish that this is a genuine form of synaesthesia, insofar as these individuals can be distinguished from nonsynaesthetes given an identical task. Tests of consistency (e.g., Baron-Cohen et al., 1993) and automaticity (e.g., Mills et al., 1999) have yet to be applied to sound-colour synaesthesia, despite the phenomenon being anecdotally reported many times before. Whilst some studies have reported their sound-colour synaesthetes to be consistent over long periods of time (e.g., Ginsberg, 1923; Ostwald, 1964), these studies have not compared performance to a control group of participants. This could be particularly important for sound-colour mappings because these are not necessarily arbitrary. For example, most people tend to associate high pitch sounds with light colours (Hubbard, 1996; Marks, 1974, 1982, 1987; Melara, 1989) and most people are able to generate, on demand, visual imagery to music (e.g., Karwoski and Odbert, 1938).

Whilst there is mileage in any study showing that a given type of synaesthesia is genuine, research must also focus upon understanding in more detail the nature of the underlying cognitive and neural processes. In particular, research should consider how synaesthesia can be used to inform theories of normal cognition. Sound-colour synaesthesia is an attractive candidate to address this question because there is some evidence to suggest that this variety of synaesthesia might be present in us all from a young age (Lewkowicz and Turkewitz, 1980; Marks et al., 1987). Moreover, cross-modal audiovisual mechanisms exist in the normal population more generally (e.g., Calvert, 2001).

Neuroanatomical studies of young kittens show pathways between auditory and visual brain centres (Dehay et al., 1984), and human neonates show visual evoked potentials given auditory stimuli (e.g., Maurer, 1997). These observations have lead Baron-Cohen and colleagues to propose a neonatal pruning hypothesis of synaesthesia (Baron-Cohen, 1996; Baron-Cohen et al., 1987, 1993, 1996). They assume that everyone is born with synaesthesia-like neural hardware but that most of us lose this during the natural course of development. Synaesthetes, perhaps because of a genetic modification, do not lose this neural hardware and retain synaesthesia into adulthood. As such, this account proposes that there are specialised pathways in adult synaesthetes that are not present in most other individuals. Thus, in this account, synaesthesia is not an extension of the cross-modal mechanisms common to all adults.



Fig. 1 – Sound-colour synaesthesia could plausibly arise from (I) direct connections between unimodal auditory and visual centers that are not normally found in other people, or from (II) feedback from multi-modal to unimodal regions that resembles normal cross-modal perception.

An alternative viewpoint would be that soundcolour synaesthesia utilises pathways that are used to integrate visual and auditory information as part of the normal mechanisms of cross-modal perception. There are cross-modal audiovisual areas of the brain that appear to respond more to the combined presence of vision and sound than to either, or the sum of, vision or sound alone (e.g., Calvert, 2001). Behaviourally, this may be important for tasks such as lip-reading (e.g., Calvert, 2001) or attentional orienting mechanisms (e.g., Spence and Driver, 1997). If there is cross-modal 'disinhibited feedback' from audiovisual areas of the brain to unimodal visual areas of the brain then this could produce overt visual responses to auditory stimuli (e.g., Grossenbacher, 1997; Grossenbacher and Lovelace, 2001). Given that these pathways are common to us all, synaesthesia would be seen as an extension or exaggeration of these particular pathways rather than requiring an extra pathway. Figure 1 simplistically compares the feedback account and the direct pathway account. Disinhibited feedback could also account for the fact that sound-colour synaesthesia can be acquired by damage to the retino-cortical pathway (e.g., Armel and Ramachandran, 1999; Vike et al., 1984) or from ingestion of hallucinogenic drugs (e.g., Hartman and Hollister, 1963). The latter are thought to operate by impeding the transmission of information from the lateral geniculate nucleus to the cortex (Hollister, 1968). Thus, both forms of acquired sound-colour synaesthesia can be viewed as a compensatory 'switching on' of visual cortex from non-visual information as a result of visual

sensory deprivation. Given that pharamcological synaesthesia can be induced within an hour it is reasonable to assume that it uses pre-existing pathways in the brain (Hollister, 1968), although it remains to be seen whether the characteristics are similar to those exhibited in the developmental form.

Further evidence to suggest that synaesthesia could be an extension of normal cross-modal perceptual mechanisms comes from the observation that both synaesthetes and non-synaesthetes match colours with sounds in a non-arbitrary way. The most commonly studied domain is pitch¹. In the synaesthesia literature, it appears that almost all cases reported to date show the trend that high pitched sounds are lighter/brighter in colour and low pitched sounds are darker in colour (e.g., Cutsforth, 1925; Riggs and Karwoski, 1934; Whitchurch, 1922; Zigler, 1930). However, it is to be noted that the authenticity of these historical cases cannot be verified. The same trend is found in non-synaesthetes. This is found if people are asked to select colours for a pitch (Marks, 1974), rate how well colours and pitches go together (Hubbard, 1996), or even in Stroop-like tasks in which participants make a light/dark judgement whilst listening to an irrelevant high or low pitched tone (Marks, 1987; Melara, 1989). It is tempting to conclude that the synaesthetes and nonsynaesthetes are using the same cognitive and neural mechanisms to derive these associations, but

¹ By which we mean pitch height, rather than pitch class (e.g., C, D#, Bb). We return to this distinction in the Discussion.

with only the former eliciting an overt, definite perceptual experience. However, to date no study has made a direct comparison between soundcolour synaesthtes and controls using the same task.

If several different notes are played simultaneously, as in a chord, most synaesthetes report experiencing several colours rather than a fusion of colours (Cutsforth, 1925; Ginsberg, 1923; Myers, 1911; Riggs and Karwoski, 1934; Rogers, 1987; Whipple, 1900). This suggests that the critical auditory unit is at the level of the single note (or pitch) rather than a unit of perception operating across several notes (e.g., pitch groupings, interval values or melody).

Synaesthesia involving aspects of music other than pitch have been documented in far less detail. Other potential influences would include timbre (related to the instrument), loudness, tempo, emotional response, consonance/dissonance, the name or pitch class of the note (e.g., C, F#) and the key or tonality of a piece (e.g., C major) (e.g., see Marks, 1975). Timbre is the aspect of sound that enables a perceiver to judge that two sounds are different when they are otherwise matched for pitch and loudness (e.g., to distinguish between the same notes on a trumpet and cello). Zigler (1930) reported two cases of synaesthesia that experienced both colours and shapes in response to tones. Whereas pitch influenced the shade, the timbre had more effect on the shape: "Every instrument excites a specific form, which maintains roughly the same features at all pitches, intensities and durations" (p. 285). In other cases it has been reported that both pitch and timbre affects the (Cutsforth, 1925: colour Masson. 1952: Whitchurch, 1922) although in some cases the same colour is elicited by a given pitch irrespective of timbre (Myers, 1911). Mudge (1920) notes that non-synaesthetes also tend to ascribe different colours to different timbres (e.g., flute and clarinet are light, whereas trombone and bass drum are dark). However, this study conflated timbre with the different pitch registers of the instruments.

In this study, Experiment 1 compares the influence of pitch and timbre on synaesthetic colour experiences. We also consider, in less detail, differences between pairs of notes (dyads) and single notes. By comparing the performance of synaesthetes directly to controls, we hope to shed light on the question of whether this form of synaesthesia may be using some of the same mechanisms that are available to us all. Experiments 2 and 3 are concerned with the automaticity of cross-modally induced synaesthesia.

EXPERIMENT 1: SOUND-COLOUR ASSOCIATIONS

The aim of this experiment is to demonstrate that sound-colour synaesthetes show a higher

internal consistency than control participants, and also to investigate how different aspects of an auditory stimulus are mapped on to the visual domain in both synaesthetes and controls.

Method

Participants

Ten sound-colour synaesthetes and 10 control participants took part. The 10 control participants consisted of 8 females and 2 males and none of them reported any synaesthetic experiences. The 10 synaesthetes had a mean age of 42.5 years, and consisted of 6 females (DSL, LHM, MH, SAW, SAWE, KZ) and 4 males (JW, JH, SJT, RJ). All ten synaesthetes reported grapheme-to-colour synaesthesia in addition to sound-to-colour synaesthesia. Over a 2-3 month interval their colour association to written letters (N = 26) and numerals (N = 10) was 86% consistent. This compares to a 33% level of consistency from 43 control subjects given a retest after only 2 weeks [t (51) = 9.31, p < .001]. Three of our sample (JW, LHM, KZ) also report vivid complex shapes in response to tones (e.g., Zigler, 1930) upon which colours are superimposed. Other synaesthetes do not report shapes as such, although the colour responses are typically accompanied by a sense of 'space' (e.g., in terms of location and/or size). For example, on hearing a siren outside of the window, one subject remarked that it was like a thin bright line moving across his visual field, similar to the trail left when someone waves a sparkler. These aspects of their sound-induced synaesthesia are not examined in this study and we concentrate solely on colour. None of the synaesthetes reported having absolute pitch (i.e., the ability to identify or produce a particular note without an external reference note) but we turn to this again in the discussion.

Sound Stimuli

Seventy different sound stimuli were used. The sounds were all of 3 seconds duration recorded as 16-bit stereo files, at a sampling frequency of 44,100 Hz (CD quality) and 65 dB. The sounds were selected to represent a range of pitches and timbres, and also to contrast single notes (N = 40)with pairs of simultaneous notes or dyads (N =30). The 40 single notes consisted of 10 piano notes, 10 pure tones (sine wave), 10 string notes, and 10 further notes taken from 10 different instruments (e.g., flute, bagpipe). The piano, string and pure tone notes were matched for frequency of the fundamental and consisted of the notes from C1 (33 Hz) up to Eb6 (1245 Hz) separated by intervals of musical fifths (i.e., 700 cents). Thus, the following notes were used: C1, G1, D2, A2, E3, B3, F#4, Db5, Ab5 and Eb6.

The remaining 10 single notes from the different instruments were all matched for fundamental frequency corresponding to that of middle C (i.e., 262 Hz). The 30 dyads also consisted of 10 piano dyads, 10 pure tone dyads and 10 string dyads. The dyads comprised of simultaneous pairs of notes separated by an interval of a musical fifth ranging from C1 (33 Hz) up to Bb6 (1865 Hz). Thus, the following dyads were used: C1 + G1, G1 + D2, D2 + A2, A2 + E3, E3 + B3, B3 + F#4, F#4 + Db5, Db5 + Ab5, Ab5 + Eb6, Eb6 + Bb6. In this way, we are able to contrast the colours assigned to isolated notes with the same notes presented in pairs.

Colour Selection

The experiment used the standard 'Windows API Choose Color' dialog box that was summoned by the Visual Basic routine used to run the test. This was chosen because it is likely to be familiar to most regular computer uses and has previously been used in studies of synaesthesia (e.g., Odgaard et al., 1999). It offers the participants a wide array of colours to select from. Colours could be chosen from one of two sets: either a basic palette of 48 preset colours or a more fine-grained colour specification made by moving a cross-hair cursor over a colour matrix. Participants could customise the colour further by adjusting a vertical slider that controls luminance. Given that synaesthetes report very precise colour experiences, we predict that they would be more likely to use the finer-grained customised colours than control participants. Each colour selection can be represented numerically by a single RGB (red, green, blue) vector with values ranging from 0 to 255 (where 0,0,0 is black and 255,255,255 is white). It is possible to convert RGB values to an approximate position in Munsell space using look-up tables used in software-based systems for displaying Munsell (Campadelli et al., $(2000)^2$. The Munsell colour space is considered a more psychophysically plausible representation of colour in which colour is described according to three dimensions: value (light to dark on a 0-10 scale), chroma (grey to coloured) and hue (a circularly varying dimension, red \rightarrow yellow \rightarrow green \rightarrow blue \rightarrow purple \rightarrow red). Standard RGB colours are a subset of the Munsell colours, but some Munsell colours cannot be represented in RGB space (and thus could not be selected). This will be a source of noise in our data. However, given that synaesthetes, but not controls, are putatively matching to a perceptual experience we might expect it to disadvantage them more (i.e., the method 'stacks the deck' against our hypothesis).

Procedure

The 70 sound stimuli were presented in a random order and the test was self-paced. Participants were seated in front of a computer in a dimmed room and put on a pair of Sennheiser HD250 headphones. They were given instructions and practice at using the colour picker before the start of the experiment. The screen contained a 'Play' button which played the sound and a 'Colour' button to call up the colour picker. Participants could play the sound more than once if they wished and go back and modify the colour. However, after participants had opted to move on to the next trial they were unable to go back and change a previous colour or listen to the old sound. It took between 20 and 40 minutes to complete all 70 selections.

The whole procedure was repeated on two further occasions. The first repetition took place within the same session after a 10 minutes interval, using a different randomised order. The second repetition took place after a 2-3 month interval.

Results

Internal Consistency over Time

The consistency of the colours generated for each sound was measured quantitatively by calculating the vector difference in RGB choices between test and retest. For example, if an RGB value of R = 255, G = 200, B = 50 was selected at Time 1 and R = 255, G = 250, B = 75 was selected at Time 2 then the distance between these points in RGB space is given by $\sqrt{(0^2 + 50^2 + 25^2)} = 55.9$. These results were analysed using a 2 × 2 ANOVA on the RGB distance scores. The two factors were group (synaesthete, non-synaesthete) and retest interval (same day, 2-3 months). The results are summarised in Figure 2.

There was a main effect of group, indicating that the synaesthetes were significantly more consistent over time than controls [F(1, 18)] =22.31, p < .01]. There was no main effect of time although it was borderline significant [F (1, 18) = 3.49, p = .08] and no interaction between them [F $(1,1 \ 8) = .01, n.s.$]. Thus, the tendency for synaesthetes to be more consistent in their colour associations over time than controls holds for both short and long retest intervals. As far as we are aware, this is the first time that this result has been reported for this variety of synaesthesia. On the same day retest, only one synaesthete (MH) out of ten fell within the control range of scores and at the 2 month retest, two synaesthetes (MH and RJ) fell within the control range of scores.

Figure 3 illustrates the colours selected for the 10 single piano, 10 single string and 10 single sine waves on two occasions, for two of our subjects (the synaesthete LHM and a control participant). It is to be noted that both show a trend for higher pitch notes to elicit lighter colours although LHM

 $^{^2}$ Note, a published Munsell colour chart was also available to ensure that the monitor was adequately calibrated.



Fig. 2 – Consistency over time, following a same day retest and a subsequent retest several months later (SD for synaesthetes is 26 and 36 resepctively, and for controls 26 and 23).

is clearly more internally consistent in her assignment of colours to tones. These are typical colour responses to this task (for information, LHM's RGB consistency, as calculated above, was 55.8 whereas CE's value was 115.0).

The same results are obtained if RGB space is transformed into separate Munsell dimensions. Treating differences in hue, chroma and value as dependent variables in 2×2 ANOVAs (group X time) shows significant effects of group for hue [F (1, 18) = 31.38, p < 001], value [F (1, 18) = 18.37, p < 001] and chroma [F (1, 18) = 24.63, p < 001]. That is, synaesthetes were significantly more consistent on all three dimensions than controls. There were main effects of time for value [F (1, 18) = 7.04 p < .05] but not hue [F (1, 18) = 1.33, n.s.] or chroma [F (1, 18) = 3.76 p = .07], and there were no interactions between group and time (F < 1).

Specificity of Colours Selected

Colours could be selected either from a preset palette of 48 different colours or the participant could produce a colour of their own by modifying one of the preset colours. It is possible to calculate how many preset colours each person chose versus how many customised colours they chose. Our hypothesis is that synaesthetes will experience very specific colours in response to tones (e.g., a particular yellow, rather than a generic yellow) and that this would be reflected in their use of the customised colour picker.

The results show that this hypothesis was borne out. Synaesthetes chose significantly more customised colours (80.7%, SD = 22) than the control group [61.6%, SD = 19.2; t (18) = 2.07, p < .05]. Thus, not only are synaesthetes more consistent over time in their choice of colours, but they also tend to choose very specific colours to represent their experiences.

Effects of Pitch and Timbre on Lightness and Chroma

It was noted in the Introduction that previous studies have shown that both synaesthetes and nonsynaesthetes display a tendency to associate low pitch notes with dark colours and high pitch notes with light colours. However, to the best of our knowledge, a direct comparison of synaesthetes and non-synaesthetes has never been conducted using the same task and stimuli. In this section, we consider single note stimuli. Dyads will be considered in a subsequent section and in somewhat less detail.

The Munsell <u>lightness</u> values were analysed in a $10 \times 2 \times 3$ ANOVA contrasting pitch ($\times 10$), group (synaesthete, non-synaesthete) and timbre (pure tone, piano, string), averaging across the different testing sessions. Both groups show a monotonic increase in lightness with pitch, suggesting a



Fig. 3 – An example of the colours selected (on 2 occasions) for the 10 single piano, sine (or pure tone) and string notes for a synaesthete (LHM, top) and a control subject (CE, bottom).

common underlying mechanism. This is illustrated in Figure 4. There was no effect of timbre and timbre did not interact with any other factors (all F's < 2), and so this dimension is not graphically displayed. There was a highly significant effect of pitch [F (9, 162) = 56.32, p < .001], but no difference between synaesthetes and controls [F (1, 18) = 2.51, n.s.] although a significant group X pitch interaction [F (9, 162) = 3.62, p < .05] reflects a tendency for synaesthetes to select somewhat lighter colours at the highest pitches.

The same analysis was conducted on the Munsell <u>chroma</u> values, although we did not have specific predictions about this dimension³. The

 $^{^3}$ Note, hue is a circularly varying dimension and cannot be analysed in the same way.



Fig. 4 – The relationship between pitch and lightness for synaesthetes and controls for the 30 single notes (collapsed across timbres and testing sessions): 0 = darkest, 10 = lightest.

chroma values were analysed in a $10 \times 2 \times 3$ ANOVA contrasting pitch (\times 10), group (synaesthete, non-synaesthete) and timbre (pure tone, piano, string), and averaging across the different testing sessions. Pitch exerts an effect on chroma [F (9, 162) = 9.28, p < .001] and not just lightness. However, chroma differs from lightness in two key respects. Firstly, there is no monotonic relationship. Instead, chroma peaks at mid-range pitches (in our experiment, a semitone below 'middle C'). Secondly, there was a highly significant main effect of timbre on chroma [F (2, 36) = 16.89, p < .001] that was not observed for lightness. Thus, musical notes from the piano and strings are, literally, more colourful than pure tones. This is illustrated in Figure 5. Importantly, there was no group difference between synaesthetes and controls [F (1, 18) = 1.26, n.s.] and none of the interactions approached significance. In sum, different aspects of auditory stimuli (pitch, timbre) appear to map on to different aspects of colour (lightness, colourfulness) in systematic ways.

However, these regularities appear to be common to synaesthetes and control subjects alike.

Finally, we investigated differences in lightness and chroma across the 10 tones taken from different instruments (bagpipes, flute, etc.) that were matched for pitch of the fundamental frequency (at middle C). Given the results presented above, one might expect differences to be manifest on the chroma but not lightness dimension. Two separate 2×10 ANOVA's were conducted with group (synaesthete, control) and instrument (N = 10) as independent variables and chroma and lightness as dependent variables. Contrary to our prediction, there were main effects of timbre on both lightness [F (9, 162) = 11.56, p < .001] and chroma values [F (9, 162) = 3.29, p < .001]. Yet again, however, there was no difference between synaesthete and control groups and no interactions between group and timbre (all F's < 2). This is illustrated in Figure 6. Even though the frequency of the fundamental was matched, it is quite conceivable that differences in perceived



Fig. 5 – The relationship between pitch, timbre and chroma for synaesthetes (top) and controls (bottom) presented with single notes. The notes are most colourful around the B3 pitch (a semitone below 'middle C') and for musical notes (string, piano) relative to pure tones.

pitch (and hence lightness) occurred across the instruments. This is because the perceived pitch is derived from the wider spectral properties of the tone and not just the fundamental frequency (e.g., Plomp, 1967). As such, a more convincing demonstration of different visual correlates of pitch and timbre is required. However, the main point that we wish to emphasise is that controls and synaesthetes appear to map auditory properties to visual properties in very similar ways.

Single Tones versus Dyads: Single Colour or Multi-coloured Experiences?

On the 2-3 month retest trial, synaesthetes were asked whether each of the 70 tones produced a



Fig. 6 – Different instruments are associated with different amounts of lightness and chroma, but these differences are found equally in synaesthetes and controls.

single colour or several different colours. This was motivated both by the initial verbal reports of our synaesthetes and also from similar observations made in the historical literature on synaesthesia (Cutsforth, 1925; Ginsberg, 1923; Myers, 1911; Riggs and Karwoski, 1934; Rogers, 1987; Whipple, 1900). Given that none of our controls reported actual colour experiences, we did not consider it pertinent for them to be asked about this. The results are summarised in Figure 7. A 3×2 ANOVA was conducted on the type of instrument (piano, sine, string) and the type of tone (single note, dyad). There was a main effect of type of tone [F (1, 9) = 14.09, p < .005] with the dyads tending to elicit two or more colours. There was also a main effect of type of instrument [F (2, 8) = 7.28, p <.05], with piano and strings being more likely to elicit two or more colours than pure tones.



Fig. 7 – Dyads and timbres with multiple harmonics often elicit more then one colour.

This could either be attributable to an effect of timbre per se or due, more simply, to the fact that natural musical tones have more complex spectral properties. There was no significant interaction [F (2, 8) = 1.59, n.s.].

The dyads also show a monotonic increase in Munsell lightness with increasing pitch (collapsing across timbres). A 2×10 ANOVA of note type (dyad vs. single note) and pitch (10 levels) on the synaesthetes revealed differences in lightness across pitch [F (9, 81) = 56.29, p < .001] but no difference between dyads and single notes [F(1, 9)]= .40, n.s.] and no interaction [F (9, 81) = .63,n.s.]. Thus, the lightness assigned to single notes (e.g., C1) is also found when these notes are the bass notes of a dyad (e.g., C1 + G1) - at least for intervals of a fifth. Control subjects show the same pattern: a significant effect of pitch [F (9, 81) =20.84, p < .001], no difference between dyads and single notes [F (1, 9) = 1.78, n.s.] and no interaction [F (9, 81) = 1.74, n.s.].

Given that participants were only allowed to choose one colour for each tone, one may wonder how they made their judgement in the case of dyads that tended to elicit two colours. An inspection of the colours reveals that (with the exception of JW) the colour of the dyads typically resembles the colour of the component notes that make it up. Given that the dyads consisted of notes of similar pitch height and always of the same timbre, the colours tended to be quite similar to each other (e.g., two slightly different shades of brown). There is at least one exception to this that is noteworthy. JW tends to report single colours to dyads and he appears to use a different mechanism to derive the colour of dyads than that used to derive the colour of isolated notes. For example, a 370 Hz single piano note elicits a darkish offyellow colour and a 554 Hz piano note elicits a slightly lighter, purer yellow colour (reliably so over time). However, when the 370 Hz and 554 Hz notes are played simultaneously as a dyad then a pinkish colour is consistently reported that bears no obvious relationship to the colour of the isolated notes. It suggests that the unit of representation that drives the association of colours in JW operates over several notes (e.g., the interval) whereas for most other synaesthetes it operates at the single note level.

Summary

In summary, Experiment 1 has demonstrated that sound-colour synaesthesia can objectively be distinguished from sound-colour associations generated by the general population in terms of internal consistency and specificity of colours selected. As such, it resembles previous reports on grapheme-colour synaesthesia (e.g., Baron-Cohen et al., 1993) and synaesthesia in which taste is the elicited experience (Ward and Simner, 2003). However, we have been able to take our analysis one step further by considering the possible cognitive mechanisms employed by both synaesthetes and controls. Both synaesthetes and controls appear to employ an *identical* strategy for

mapping sounds to colours that involves pitchlightness mappings, and they show identical influences of timbre on chroma. We therefore believe that this variety of synaesthesia can best be explained as an exaggeration of cross-modal mechanisms common to us all rather than a privileged, direct pathway between audition and vision that is present in synaesthetes but not in others. In synaesthetes we consider the mapping to be precise (as shown by high internal consistency) and results in a vivid perceptual experience, whereas in controls the mapping is noisy and is not associated with reports of actual visual experience. One further difference between the two groups may lie in the automaticity of the mapping between sound and colour. To test this, the next experiment uses a cross-modal variant of the synaesthetic Stroop effect.

EXPERIMENT 2: CROSS-MODAL STROOP INTERFERENCE

The aim of this experiment was to establish that, in the synaesthetes but not controls, the colour experiences are automatically elicited upon presentation of a tone. This is analogous to the Stroop-like interference effect found in graphemecolour synaesthesia (e.g., Mills et al., 1999). In this instance, synaesthetes and controls were asked to name the colour of a patch on the screen whilst simultaneously listening to tones over the headphones which they were asked to ignore. For the synaesthetes, the tones either elicited a synaesthetic colour that was congruent with the colour on the screen or a synaesthetic colour that was incongruent with it.

Method

Design

The experiment uses a mixed $2 \times 2 \times 2$ design with tone-colour pairing (congruent, incongruent) as one factor, sound-colour onset asynchrony as a second factor (sound precedes colour by 0 or 150 msec) and presence or absence of synaesthesia as the between-subject factor.

Participants

Nine synaesthetes from Experiment 1 were retested together with nine controls. One of the synaesthetes who took part in Experiment 1 was unavailable for further testing.

Materials

For each synaesthete, six tones were selected from the set used in Experiment 1 on the basis that they had previously elicited a consistent colour. An effort was made to select a range of different colours and, hence, a range of different pitches and timbres were also used. Incongruent colour-tone pairs were created by pairing up the 6 colours and the 6 tones in a different way. There were 120 trials in total, with equal numbers of congruent and incongruent pairings. For half of the trials the onset of the sound and the colour patch were simultaneous, and for the remainder of trials the sound preceded the colour patch by 150 msec. Each control subject was given one of the soundcolour sets of a corresponding synaesthete.

Procedure

Stimuli were presented using E-prime software. Subjects were seated in front of a computer, wearing Sennheiser headphones, and with a microphone placed close to their mouth. They were told that they would see some coloured patches on the screen and, at the same time, hear some sounds presented over the headphones. They were instructed to name the colour of the patch on the screen as quickly and accurately as possible and that they should try to ignore the sound. They were asked to produce a single colour term (e.g., "blue" rather than "dark blue" or "royal blue"). Each trial proceeded as follows. There was a central fixation cross presented for 1000 msec. Following this, the tone was played through the headphones. A coloured rectangle appeared in the centre of the screen against a mid-gray background either at the same time as the tone or 150 msec later. When subjects made their response, both the colour and the tone stopped and there was a 1000 msec interval before the beginning of the next trial.

Results and Discussion

The results are summarized in Figure 8. A 2 \times 2×2 mixed ANOVA revealed a significant main effect of SOA [F (1, 16) = 25.72, p < .001], a significant effect of stimulus congruency [F (1, 16) = 6.28, p < .05] but the effect of group (i.e., synaesthete vs. control) did not reach significance [F(1, 16) = 3.97, p = .06]. Critically, there was a significant interaction between group and stimulus congruency [F (1, 16) = 8.56, p < .01] suggesting that the synaesthetes show a greater effect of congruency/incongruency sound-colour than controls. No other interaction approached significance (p > .20). In order to examine the group × stimulus congruency interaction further, two separate repeated measures were run on the synaesthete group and on the control group. For the synaesthetes, there was a significant effect of SOA [F (1, 8) = 13.55, p < .01] and of stimulus congruency [F (1, 8) = 8.07, p < .05] but no interaction. For the control group, a similar analysis revealed only a significant effect of SOA [F(1, 8) = 13.55, p < .01]. This suggests that the



Fig. 8 – Synaesthetes show cross-modal Stroop-like interference when incongruent sounds and colours are presented together (0 msec SOA) and when sound precedes colour (150 msec SOA). For example, synaesthetes are slowed at naming a red patch when also given a 'blue sound' (SD for synaesthetes and controls across conditions is 190 and 107 respectively).

main effect of congruency reported in the initial analysis is entirely attributable to the performance of the synaesthetes.

This is the first ever demonstration of a synaesthetic Stroop effect that is cross-modal in nature. It suggests that the synaesthetic colours are automatically elicited even when the inducing stimulus is task irrelevant and even when the inducing stimulus is in a different sensory modality. It is convincing evidence for the authenticity of this type of synaesthesia over and above the intriguing commonalities that exist between synaesthetes and controls in terms of the way that they choose colours given sounds.

In a previous pilot study, one of our synaesthetes (LHM) commented that when the sound and the colour were congruent it felt as if the sound was coming out of the colour on the screen, rather than through the headphones. That is, she reported an illusory capture of the location of sound by a congruent visual stimulus. This constitutes a synaesthetic version of the 'ventriloquist effect' that has been documented in studies of normal cross-modal perception (e.g., Driver, 1996; Spence and Driver, 2000). The synaesthete had not studied psychology and was unlikely to be aware of the literature on this subject. In the current study we took the opportunity to question all the other synaesthetes

after the Stroop experiment (showing them the stimuli again if need be) but we found only two other synaesthetes who reported the phenomenon. Thus, this illusion is by no means universal. Notwithstanding this, it prompted us to ask whether non-spatial auditory cues could lead to 'covert' spatial orienting to a congruent visual stimulus even when the cue was completely uninformative.

EXPERIMENT 3: CROSS-MODAL ORIENTING OF SPATIAL ATTENTION

In the Stroop experiment, above, participants were required to ignore the sound whilst attending to the colour dimension. In the present experiment, both the sound <u>and</u> the colour are completely irrelevant and uninformative for the task at hand. The task is a cross-modal variant of the Posner cueing paradigm (Posner, 1980; Posner and Cohen, 1984). In the original task, a peripheral lateralised visual cue preceded the presentation of a lateralised visual target on either the same side or the opposite side as the cue. The cue was found to facilitate same-side detection times (up to around 200 msec SOA) even though the cue itself was uninformative (e.g., on half the trials it appeared on the wrong side). At longer delays (> 300 msec) inhibitory

effects were found associated with a cost in shifting attention back to a previously empty location. Similar effects have been found for lateralised auditory cues and visual targets (e.g., Spence and Driver, 1997). In the present experiment, the auditory cue is presented through both ears of the headphones and, as such, has no lateralised component. The auditory cue is, however, presented at the same time as two coloured rectangles on the left and right of the screen, one of which is synaesthetically congruent with the sound. The task is to detect a target (an asterisk) presented in one of the two rectangles a short time after the cue (50, 150 or 300 msec). Will an auditory cue orient attention to the synaesthetically congruent location?

Method

Design

The experiment uses a $2 \times 3 \times 2$ mixed design with location of the target as one factor (either on the congruent or incongruent colour), the onset asynchrony between auditory cue and visual target as another factor (50 msec, 150 msec, 300 msec) and the type of subject (synaesthete or control) as the final factor.

Participants

The same synaesthete and control participants as in Experiment 2 took part in this experiment.

Materials

The 6 consistent colour-tone pairings used in Experiment 2 were used again in the present experiment. For each tone, there were 20 possible permutations of colour and target location pairings - the congruent colour was either on the left or right (x2), the incongruent colour was one of the remaining colours (x5) and the target either appeared on the congruent or incongruent colour (x2). Thus, for the 6 tones there were 120 unique pairings of colour, sound and target position. With the addition of the onset-asynchrony variable between the sound and the target each experimental block consisted of 360 trials. In addition there were 90 catch trials in which the target was absent (some 20% of trials). The 450 trials were presented in a random order.

Procedure

The participants were seated 70 cm away from a colour computer monitor. They were told that they would first of all see a fixation cross (+) in the centre of the screen that they should attend to, and that they would then see two coloured squares on either side of the screen. One of the coloured squares would contain an asterisk (*) and the task of the participants was to identify the target as quickly as possible. None of the synaesthetes reported an experience of colour to the asterisk. Participants made their response by pressing the space bar upon target detection. Participants were informed that they would also hear a sound through the headphones whilst performing the task. They were told that the tone is not relevant to the task and that they should ignore it. The fixation cross was displayed for 1000 msec on a mid-grey background. Following this the coloured squares were displayed on the screen against the same midgrey background. The squares were approximately 5 cm in size and their centres were separated by 22 cm on the screen (i.e., the targets were 8.9 degrees from fixation). The auditory tone was always played at the same time as the coloured patches appeared. The target appeared after a delay of either 50, 150 or 300 msec. Both the target and the tone remained until the participant made a response or until a 1500 msec time-out (whichever came first). The next trial then began immediately.

Results and Discussion

The results are summarized in Figure 9. A 2×3 \times 2 mixed ANOVA on the reaction times revealed a significant congruency X group interaction [F(1, 16)]= 4.30, p < .05] which suggests that synaesthetic congruency between sound and colour can result in spatial orienting of attention. There was a marginally insignificant 3 way interaction between congruency \times group \times SOA [F (2, 32) = 3.11, p = .06] and a marginally insignificant effect of group [F(1, 16) =3.82, p = .07] consistent with the observation that synaesthetes tend to be slower particularly when cued to the wrong location. Other effects did not approach significance (p > .10). There was no evidence of an interaction between congruency and SOA. We would have predicted this interaction based on the standard Posner cueing paradigm. If the target does not appear in the cued location, attention is shifted to the other location and there is a cost of returning to the previous location ('inhibition of return'). Presumably it takes some (unknown) amount of time for the synaesthesia to be elicited and so it is perhaps not too surprising that there was no inhibition of return in the time window used, but we would predict it at longer SOAs. The proportion of missed targets and false alarms was low (misses = 1.06%, false alarms = 3.70%) and did not differ between synaesthetes and controls [misses t (16) = .67, n.s.; false alarms t (16) = .23, n.s.].

To summarise again, when the target appears in an incongruently coloured location with respect to the sound then the synaesthetes, but not the controls, tend to be slowed. This presumably occurs because attention is 'captured' by the congruent location. These results add further weight to the conclusion that synaesthesia is automatically elicited. It also



Fig. 9 – A non-lateralised and uninformative sound can orient attention to a synaesthetically congruent colour patch in a Posnertype experiment (SD for synaesthetes and controls across conditions is 141 and 85 respectively).

suggests that synaesthesia appears to use the same attentional mechanisms as those found in veridical perception (or an analogous mechanism that operates on the same lines).

GENERAL DISCUSSION

In this study we were able to find clear differences between sound-colour synaesthetes and controls given the same task. Differences were found in terms of internal consistency of colour associations, specificity of the colour selection, and in terms of the automaticity of colour association as measured by Stroop-like interference with incongruent sound-colour pairs and attentional orienting to a congruent sound-colour pair. This is the first time that these objective tests have been conducted on this variety of synaesthesia, although our results are entirely consistent with previous anecdotal reports in the literature.

Although we were able to find clear differences between the synaesthetes and controls on these tests, in other ways the two groups were remarkably similar. Table I summarises what we consider to be the main similarities and differences between synaesthetes and controls. First of all, we believe that both groups use the same cognitive mechanism to map between the auditory and visual domain. This mechanism appears to involve mappings between pitch and lightness. It is conceivable that different aspects of complex sounds became associated with different aspects of vision. For example, the presence or absence of timbre influences the number of colours and the richness of the colours (in terms of the Munsell chroma) observed. Whilst this remains to be fully explored there is no evidence, as yet, to suggest fundamental differences between the way that synaesthetes and others (presumably engaged in imagery) associated sound and colour. We speculate that the same cognitive mechanism is used by both groups but that the mechanism differs in terms of its precision and automaticity between the groups. For synaesthetes the mapping between sounds and colour is very precise as is evidenced by their high internal consistency and selection of customised colours. In contrast, the mapping between sound

TABLE I

Similarities and differences in sound-colour associations between synaesthetes and non-synaestetes

	Synaesthetes	General Population	
Nature of mapping	Pitch → lightness	Pitch → lightness	
Specificity of mapping	Precise	Noisy	
Mode of induction	Automatic	Strategic	
Perceptual experience	Conscious, explicit	None or implicit	

and colour in non-synaesthetes is noisy. Although controls show the same central tendency to map certain pitches on to certain lightnesss they show no strong commitment to a precise colour. Moreover, we suggest that whereas the synaesthetes evoke this mechanism automatically, the controls may evoke the mechanism only strategically. Experiment 1 demonstrated that controls can generate broadly similar sound-colour associations to synaesthetes, but Experiments 2 and 3 established that there is no evidence that controls do so automatically⁴. We speculate that this may reflect the amount of feedback connections in a model such as that in Figure 1 above.

It is important to note that whilst we consider objective differences between synaesthetes and controls to be crucial, this does not make phenomenology redundant. For example, one of our control subjects had an internal consistency comparable to synaesthetes but did not claim to have synaesthetic experiences (unfortunately, she wasn't tested on Stroop) and we would not wish to diagnose her as having synaesthesia. In our view, reported experiences together with objective tests comprise a primary source of data that need to be explained. As such, an explanation in terms of synaesthesia is more parsimonious than an explanation in terms of learned verbal associations between colour and sound, because the latter but not the former requires an ad hoc account along the lines of a highly specific delusion pertaining to their own experiences⁵. We do not wish to be glib about the importance of ascertaining that synaesthetic experiences are perceptual. This is far from easy (e.g., imagery may use perceptual resources), although our study on cross-modal synaesthetic orienting offers support for a perceptual account. However, we do wish to note the importance of considering phenomenological reports alongside objective tests. This situation is not unique to synaesthesia. For example, the interest in blindsight comes from the self-reported lack of vision together with the residual objective abilities, and not from performance on the objective tests alone (e.g., Cowey, 2004). Similarly, the search for the neural correlates of consciousness is predicated on the assumption that conscious reports can serve as primary data that require explanation (e.g., Frith et al., 1999).

As noted earlier, there are two different accounts of how synaesthetic experiences might arise. First of all, they may reflect direct hard-wired connections between unimodal auditory and unimodal visual centres that are found in synaesthetes but that are not present in other people (except perhaps in the neonatal period, e.g., Baron-Cohen, 1996; Baron-Cohen et al., 1993; Harrison, 2001). Secondly, they may reflect feedback pathways from bimodal (or multimodal) audiovisual regions to unimodal visual regions of the brain (e.g., Grossenbacher and Lovelace, 2001). These pathways are assumed to be present in us all to some extent (e.g., Calvert et al., 1999; Driver and Spence, 2000; Macaluso et al., 2000). We suggest that the fact that nonsynaesthetes appear to use the same heuristic for matching sounds to colours as non-synaesthetes is more consistent with the latter view than the former. It might be possible to develop some hybrid theory that was intermediate between these positions. For example, under the direct linkage account one could assume that the pruning between auditory and visual centres is only partially achieved in nonsynaesthetes. However, in the absence of further empirical evidence we are inclined towards the more parsimonious explanation that sound-colour synaesthesia recruits the normal mechanisms of audiovisual perception and attention.

It should be noted that our results are completely neutral with regards to the nature of the underlying neural mechanism. For example, a failure of neural pruning (or apoptosis) has been put forward as one candidate mechanism by proponents of the direct linkage theory (Baron-Cohen, 1996; Maurer, 1997). However, the same mechanism could apply equally well if one were to assume that the pruning failure was between some cross-modal nexus and unimodal visual areas. In sound-colour synaesthesia that is pharmacologically induced it is also conceivable that the same route is used although the neural mechanism is likely to be very different. For example, the drug may stimulate release of neurotransmitters from the cross-modal route and/or impede transmission to the unimodal visual pathway (Hartman and Hollister, 1963; Hollister, 1968).

It is interesting to speculate where in the brain the hypothesised cross-modal nexus giving rise to this type of synaesthesia might lie. Functional imaging studies have consistently identified the left Superior Temporal Sulcus (STS) as being involved in the integration of auditory and visual information (for a review see Calvert, 2001). This region of the brain responds to the presence of (congruent) auditory and visual stimuli more than would be expected from either modality alone or from the sum of effects from each modality (i.e., the effects are super-additive). For example, it responds strongly when presented with congruent lip movements and heard speech (Calvert et al., 2000), when presented with congruent phonemes and written graphemes (Raij et al., 2001), and also with non-speech (white noise) and meaningless vision

⁴ One other study (Melara, 1989) noted interference when subjects were asked to make light/dark judgments when ignoring low/high pitches respectively (relative to high/low pitches respectively), and also interference when asked to make low/high pitch judgments when ignoring light/dark patches respectively (relative to dark/light respectively). This differs from our experiment in which subjects were asked to respond with a colour name rather than a lightness judgment, so the level of interference may be reduced or eliminated for controls.

⁵ It could be the case that cross-modal links do become incorporated into linguistic structures (Osgood, 1960), although the fact that cross-modal links emerge early (Lewkowicz and Turkewitz, 1980) suggests that the effects are not purely language related.

(checkerboards) when presented in temporal synchrony (Calvert et al., 2001). One obvious experiment would be to determine how this region responds, in non-synaesthetes, when presented with high/low pitch together with dark/light shades. In this instance, the congruency between vision and audition is not determined spatially (as in lip reading), temporally (as in Calvert et al., 2001) or meaningfully (as in phoneme-letters) but is determined by the stimulus dimensions themselves. It would then be interesting to determine whether this region responds strongly, in sound-colour synaesthetes, when given unimodal auditory stimuli.

In the synaesthetes reported here, the mapping from sound to colour appears to be determined by pitch height rather than pitch class. Pitch height refers to a monotonically varying dimension (related to the ascending order of octaves), whereas pitch class is a circularly varying dimension that describes position within an octave (Shepard, 1964). For example, piano notes of 220 Hz and 440 Hz have different pitch heights but the same pitch class because both notes are 'A' (or "la") on the Western musical scale. This octave equivalence is not simply a result of musical training or a cultural artefact of this system of music. It is found in many different cultures (Nettl, 1956) and may relate to a 2:1 ratio of pitch heights. Thus, the fact that 440 Hz (as opposed to, say, 450 Hz) is denoted as 'A' is purely cultural whereas the fact that notes of 220 Hz, 880 Hz and 110 Hz are considered perceptually similar to 440 Hz is assumed to reflect an underlying property of auditory processing. For some synaesthetes, pitch class appears to be the strongest determinant of elicited colour (Bernard, 1986; Carroll and Greenberg, 1961; Haack and Radocy, 1981; Langfeld, 1914; Myers, 1915; Peacock, 1985; Riggs and Karwoski, 1934; Rogers, 1987; Vernon, 1930). Thus, all notes corresponding to the pitch class of 'A' would have the same colour in spite of differences in pitch height. In some cases, both pitch class and pitch height are important. The case reported by Riggs and Karwoski (1934) assigned colours to notes such that the same note one octave apart had the same colour, but with the higher note being lighter. All the reported cases in which pitch class determines the colour of a note also have absolute pitch. Absolute pitch is the ability to identify or produce a note without an external reference (Ward, 1999). It is related to early training and possibly musical а genetic predisposition (Baharloo et al., 2000). The parallels with synaesthesia are intriguing. It has been suggested that absolute pitch could be an outcome of neural pruning (Keenan et al., 2001) and that all infants are born with the ability (Ward, 1999). It has also been claimed that many people possess a more implicit absolute pitch in which songs can be sung in their normal key despite an inability to label the note or key even after training (Levitin,

1994). It is an empirical question whether the sample of synaesthetes reported here could use their elicited colours in order to deduce the identity of a note so as to acquire a kind of absolute pitch (none of our sample report the ability, despite a few of them being musically competent). A number of conditions would be necessary for this to occur, including (a) that notes a semi-tone apart have perceivable differences in colour (b) these differences are not modified by the context in which the note is played (c) the perceptual differences can be committed to long-term memory and (d) these stored perceptual differences can acquire linguistic labels. These are open, but testable, questions. At present, all that we wish to note is that pitch is a multi-dimensional attribute and this may produce multiple forms of synaesthesia - only one of which we have been able to document in the present study.

In sum, our conclusion is that the type of sound-colour synaesthesia reported here is a genuine phenomenon in which pitch heights map on to colours as a function of the lightness. We suggest that this variety of synaesthesia recruits normal mechanisms of cross-modal perception and attention and can therefore be used to speak to theories of normal cognition. For example, it raises important questions about the relationship between strategically-generated imagery versus perception. The study of synaesthesia may also provide a way of disentangling different aspects of musical cognition (e.g., the different dimensions underlying pitch and timbre perception). Finally, it might also shed light on the necessary and sufficient neural substrates of conscious perceptual experience. These are the challenges for the future.

Acknowledgments. The Leverhulme Trust (Ref: F/07134/AG).

References

- ARMEL KC and RAMACHANDRAN VS. Acquired synaesthesia in retinitis pigmentosa. *Neurocase*, 5: 293-296, 1999.
- BAHARLOO S, SERVICE SK, RISCH N, GITSCHIER J and FREIMER NB. Familial aggregation absolute pitch. American Journal of Human Genetics, 67: 755-758, 2000.
- BARON-COHEN S. Is there a normal phase of synaesthesia in development? *Psyche*, 2: 1996.
- BARON-COHEN S, BURT L, SMITH-LAITTAN F, HARRISON and BOLTON P. Synaesthesia: Prevalence and familiality. *Perception*, 25: 1073-1079, 1996.
- BARON-COHN S, HARRISON J, GOLDSTEIN LH and WYKE M. Coloured speech perception: Is synaesthesia what happens when modularity breaks down? *Perception*, 22: 419-426, 1993.
- BARON-COHEN S, WYKE MA and BINNIE C. Hearing words and seeing colours: An experimental investigation of a case of synaesthesia. *Perception*, *16*: 761-767, 1987.
- BERNARD JW. Messiaen's synaesthesia: The correspondence between color and sound structure in his music. *Music Perception, 4:* 41-68, 1986.
- CALVERT GA. Crossmodal processing in the human brain: Insights from functional neuroimaging studies. *Cerebral Cortex, 11:* 1110-1123, 2001.
- CALVERT GA, BRAMMER MJ, BULLMORE ET, CAMBELL R, IVERSEN SD and DAVID AS. Response amplification in sensory-specific cortices during cross-modal binding. *NeuroReport*, 10: 2619-2623, 1999.
- CALVERT GA, CAMBELL R and BRAMMER MJ. Evidence from

functional magnetic resonance imaging of crossmodal binding in the human heteromodal cortex. Current Biology, 10: 649-657, 2000

- CALVERT GA, HANSEN PC, IVERSEN SD and BRAMMER MJ. Detection of audio-visual integration sites in humans by application of electrophysiological criteria to the BOLD effect. NeuroImage, 14: 427-438, 2001.
- CAMPADELLI P, SCHETTINI R and ZUFFI S. A system for the automatic selection of conspicuous color sets for qualitative data display and visual interface design. Journal of Image and Graphics, 5: 500-503, 2000.
- CARROLL JB and GREENBERG JH. Two cases of synaesthesia for color and musical tonality associated with absolute pitch. Perceptual and Motor Skills, 13: 48, 1961.
- COWEY A. Fact, artefact and myth about blindsight. Quarterly Journal of Experimental Psychology, 57A: 577-609, 2004. CUTSFORTH TD. The role of emotion in a synaesthetic subject.
- American Journal of Psychology, 36: 527-543, 1925
- CYTOWIC RE. Synaesthesia: A Union of the Senses (2nd ed.). Cambridge, MA: MIT Press, 2002.
- DEHAY C, BULLIER J and KENNEDY H. Transient projections from the fronto-parietal and temporal cortex to areas 17, 18, 19 in the kitten. Experimental Brain Research, 57: 208-212, 1984.
- DRIVER J. Enhancement of selective listening by illusory mislocation of speech sounds due to lip-reading. *Nature*, 381: 66-68, 1996.
- DRIVER J and SPENCE C. Multisensory perception: Beyond modularity and convergence. Current Biology, 10: R731-R735, 2000.
- FRITH C, PERRY R and LUMER E. The neural correlates of conscious experience: An experimental framework. Trends in Cognitive Sciences, 3: 105-114, 1999.
- GINSBERG L. A case of synaesthesia. American Journal of Psychology, 34: 582-589, 1923.
- GROSSENBACHER PG. Perception and sensory information in synaesthetic experience. In S Baron-Cohen and JE Harrison (Eds), Synaesthesia: Classic and Contemporary Readings. Oxford: Blackwell, 1997. GROSSENBACHER PG and LOVELACE CT. Mechanisms of
- synaesthesia: Cognitive and physiological constraints. Trends in Cognitive Sciences, 5: 36-41, 2001. НААСК РА and RADOCY RE. A case study of a chromesthetic.
- Journal of Research in Music Education, 29: 85-90, 1981.
- HARRISON J. Synaesthesia: The Strangest Thing. Oxford: Oxford University Press, 2001. HARTMAN AM and HOLLISTER LE. Effect of mescaline, lysergic
- acid diethylamide and psilocybin on color perception. *Psychopharmacolgia, 4:* 441-451, 1963.
- HOLLISTER LE. Chemical psychoses: LSD and related drugs. Springfield: Charles C. Thomas, 1968.
- HUBBARD TL. Synesthesia-like mappings of lightness, pitch and melodic interval. American Journal of Psychology, 109: 219-238, 1996.
- KARWOSKI TF and ODBERT HS. Color-music. Psychological Monographs, 50: 1-60, 1938.
- KEENAN JP, THANGARAJ V, HALPERN AR and SCHLAUG G. Absolute pitch and planum temporale. NeuroImage, 14: 1402-1408, 2001.
- LANGFELD HS. Note on a case of chromesthesia. Psychological Bulletin, 11: 113-114, 1914.
- LEVITIN DJ. Absolute memory for musical pitch: Evidence from the production of learned melodies. *Perception and Psychophysics*, 56: 414-423, 1994.
- LEWKIWICZ D and TURKEWITZ G. Cross-modal equivalence in infancy: Auditory-visual intensity matching. Developmental Psychology, 16: 597-607, 1980.
- MACALUSO E, FRITH CD and DRIVER J. Modulation of human visual cortex by crossmodal spatial attention. Science, 289: 1206-1208, 2000.
- MARKS L. On cross-modal similarity: Auditory-visual interactions in speeded discrimination. *Journal of Experimental Psycholog:* Human Perception and Performance, 13: 384-394, 1987
- MARKS LE. On associations of light and sound: The mediation of brightness, pitch, and loudness. American Journal of Psychology, 87: 173-188, 1974.
- MARKS LE. On coloured-hearing synaesthesia: Cross-modal

translations of sensory dimensions. Psychological Bulletin, 82: 303-331, 1975.

- MARKS LE. Bright sneezes and dark coughs, loud sunlight and soft moonlight. Journal of Experimental Psychology: Human Perception and Performance, 8: 177-193, 1982
- MARKS LE, HAMMERAL RJ and BORNSTEAIN MH. Perceiving similarity and comprehending metaphor. Monographs of the Society for Research in Child Development, 52: 1-93, 1987.
- MASSON DI. Synesthesia and sound spectra. Word, 8: 39-41, 1952.
- MAURER D. Neonatal synaesthesia: Implications for the processing of speech and faces. In S Baron-Cohen and JE Harrison (Eds), Synaesthesia: Classic and Contemporary Readings. Oxford: Blackwell, 1997.
- MELARA RD. Dimensional interaction between color and pitch. Journal of Experimental Psychology: Human Perception and Performance, 15: 69-79, 1989.
- MILLS CB, BOTELER EH and OLIVER GK. Digit synaesthesia: A case study using a stroop-t Neuropsychology, 16: 181-191, 1999. stroop-type test. Cognitive
- MUDGE EL. The common synaesthesia of music. Journal of Applied Psychology, 4: 453-456, 1920.
- MYERS C. Two cases of synesthesia. British Journal of Psychology, 7: 112-117, 1915.
- MYERS CS. A case of synaesthesia. British Journal of Psychology, 4: 228-238, 1911.
- NETTL B. *Music in primitive culture*. Cambridge: Harvard University Press, 1956.
- ODGAARD EC, FLOWERS JH and BRADMAN HL. An investigation of the cognitive and perceptual dynamics of a colour-digit synaesthete. *Perception*, 28: 651-664, 1999.
- OSGOOD CE. The cross-cultural generality of visual-verbal synesthetic tendencies. Behavioral Science, 5: 146-169, 1960.
- OSTWALD PF. Color hearing. Archives of General Psychiatry, 11: 40-47, 1964.
- PEACOCK K. Synaesthetic perception: Alexander Scriabin's color hearing. Music Perception, 2: 483-505, 1985
- PLOMP R. Pitch of complex tones. Journal of the Acoustical Society of America, 41: 1526-1533, 1967.
- POSNER MI. Orienting of attention: The VII Sir Frederic Bartlett Lecture. Quarterly Journal of Experimental Psychology, 32: 3-25, 1980.
- POSNER MI and COHEN Y. Components of visual orienting. In H Boumer and DG Bouwhuis (Eds), Attention and Performance X. Hove: Lawrence Earlbaum Associates, 1984.
- RAIJ T, UUTELA K, KRAIJ T and HARI R. Audiovisual integration of letters in the human brain. *Neuron*, 28: 617-628, 2001.
 RIGGS LA and KARWOSKI T. Synaesthesia. *British Journal of Psychology*, 25: 29-41, 1934.
- ROGERS GL. Four cases of pitch-specific chromesthesia in trained musicians with absolute pitch. Psychology of Music, 15: 198-207, 1987
- SHEPARD RN. Circularity in judgements of relative pitch. Journal of the Acoustical Society of America, 36: 2345-2353, 1964.
- SPENCE C and DRIVER J. Audiovisual links in exogenous covert spatial orienting. Perception and Psychophysics, 59: 1-22, 1997.
- SPENCE C and DRIVER J. Attracting attention to the illusory location of a sound: Reflexive crossmodal orienting and ventriloquism. NeuroReport, 11: 2057-2061, 2000.
- VERNON PE. Synaesthesia in music. Psyche, 10: 22-39, 1930.
- VIKE J, JABBARI and MAITLAND CG. Auditory-visual synesthesia: Report of a case with intact visual pathways. Archives of Neurology, 41: 680-681, 1984.
- WARD J and SIMMER J. Lexical-gustatory synaesthesia: Linguistic and conceptual factors. *Cognition*, 89: 237-261, 2003.
- WARD WD. Absolute pitch. In D Deutsch (Ed), The Psychology of
- Music (2nd ed.). London: Academic Press, 1999.
 WHIPPLE GM. Two cases of synaesthesia. American Journal of Psychology, 11: 377-404, 1900.
- WHITCHURCH AK. Synaesthesia in a child of three and a half years. American Journal of Psychology, 33: 302-303, 1922
- ZIGLER MJ. Tone shapes: A novel type of synaesthesia. Journal of General Psychology, 3: 277-286, 1930.

Jamie Ward, Department of Psychology, University College London, Gower Street, London WC1E 6BT, UK. e-mail: jamie.ward@ucl.ac.uk