Space and Sound: a two-factor approach to pitch perception

Adam Morton – adam.morton@ubc.ca

 version of September 2014

"He [Sibelius] once baffled a group of Finnish students by giving a lecture on the overtone series of a meadow." Ross 2007

abstract: I identify two components in the perception of musical pitches, which make pitch perception more like colour perception than it is usually taken to be. To back up this surprising claim I describe a programme whereby individuals can learn to identify the components in musical tones. I also claim that following this programme can affect one's pitch-recognition capacities.

contents

1. the project & the claims

I have been trying to find new ways of describing the ways musical pitches sound, and investigating whether having these descriptions feeds back to perception. A practical phenomenology (see Dennett 1991.) The project may not be that unusual, among half-competent musicians and musically obsessed philosophers, and during several years I would often think idly about it, and do not very successful exercises with tuning forks. (I became good at unconsciously estimating which fork it was from its weight!) I have devoted the summer of 2014 to the project: to making progress or realizing that the idea is hopeless. At first things were pointing to the second possibility, but then there was a breakthrough, associated with the spread/wobble distinction below. Now I find that I can classify pitches, independently of their register, into several rough overlapping classes. I often have intuitions now about what key music is in: though I am sometimes very wrong I am quite often right. And though I make mistakes, sometimes grotesque ones, I can often classify a pitch within a tone or so. (Success varies with timbre, duration, and whether the sound is coming through stereo headphones or drifting through a window.)

The method I describe has raised me from a very low level of performance at pitch and key recognition to a mediocre one[1]. That's progress! (But this does raise the question of what it can do for someone beginning from a higher level.) My results may not be different from those that others use to similar ends, for example by knowing their approximate singing range. (Though I think it is potentially more accurate.) But it is interesting. And it leads to curious and interesting experiences, in which one hears tones as having the two factors I will describe. See section **7 (v)** for a dramatic such experience; see also section **10**. In terms of these, I can begin to say how different pitches sound different, and to use the descriptions to remember them. Continuing the list, though these are all aspects of the same few realisations, I much more readily now hear sounds outside a musical context as

1 I shall describe myself as a musical incompetent. In fact, I am a curious patchwork of capacities and failings. But the details are probably irrelevant.

having a determinate pitch, even when I cannot name it. When sounds are presented as music, I often now have definite impressions of the pitch of a note or the key of a passage, even when I cannot force this into the standard language of lettered notes. (The analogy with colour below might help give a sense of this.)

Besides these practical effects there are interesting issues of perception and language, of describing sounds in novel ways. This essay describes what I have learned and what I have become able to do, in *roughly* the order I learned them. But I recommend reading through it all before trying it. If I am right, I have discovered a basic aspect of pitch perception, and a technique that allows me to exploit it. Individual history and perceptions vary, though, and you may have to experiment to apply these ideas to your own case.

I am choosing my words to prevent you from taking my claims either as trivial or as stronger than they are. An analogy with colour perception may help. Suppose that our thinking about colour was like our language and awareness of pitch. (I am not denying that there are fundamental reasons why it is not.) There would be a very small minority of super-acute chromatic sensitives who could confidently say "that fabric is just a little more blue than my mother's car, which she sold five years ago", or "if you use the colour on this sample you'll find when you get home and do a whole wall that it is just a little bit lighter than its neighbour." Most of these people would have known from an early age that they could do this. Meanwhile, the rest of us would stumble along with "bright", "vivid", and "now that I see them side by side they look very different." No doubt some people are much sharper about colour than others, but most of us are in neither of the two categories. We have vague useful words like "red", "light brown" and "mauve", of varying degrees of precision. And we can roughly compare present and remembered hues in a way that accords with the limits of our verbal precision. Children have to learn to see colours in accordance with colour language, and when they are learning the vocabulary they make mistakes that seem bizarre to adults.

We would also find, though this would be known only to experts and seem surprising to everyone, that most people represent precise hues more precisely than they can consciously describe them. If people were forced to 'guess' at colour matches of, say, their childhood clothes or the walls of rooms they visited briefly, they would do this more reliably and accurately than they would have predicted. But when asked to say in words what colours these clothes and walls were, they would perform much less well. (I predict that something like this is actually true with level of illumination and degree of saturation.) These facts would suggest that a more acute classification lurks in the perceptual apparatus, lacking a procedure to bring it to language and consciousness.

The analogy with colour perception can be taken further. We see colours by using retinal cone cells and associated neural systems sensitive to three overlapping ranges of wavelengths. (It would not work if there were only one: two can do it and does in some animals including a few humans. Ali & Klyne 1985.) The analysis and method I am describing postulates two sensitivities. They are combined to get pitch judgements, which are usually unconscious and not available to language. But each of them can be made and named separately. Then they can be combined in a process of conscious reasoning. By ordinary people! That is the central and obviously controversial claim.

Another way of making the claim. Distinguish access from accuracy. People with absolute pitch vary in how precisely they can discriminate. With many there is a definite element of categorial perception: notes are assimilated to a fixed number of labels, usually the standard twelve semitones (Deutsch 2013b). So as a limiting case imagine someone who can verbalize, consciously access, whether a note is within a half-octave centred on B. They know when it is A, A#, B, C, C#, or D and when alternatively it is D#, E, F, G, or G#, presumably with a region of uncertainty at the boundaries between the two (so D and G# would elicit less certain judgements). These two regions are essentially what I will later call the + and – categories. Such a person would be able to do things that most people cannot –

they would have conscious access to a roughly absolute pitch information – but they would not have a very accurate sense of pitches. My aim is to develop access, letting accuracy fall as it may. (A conjecture: accuracy can be refined, but access requires special gifts or an unusual approach. Given access and hard work accuracy will improve at least somewhat.)

We do have vague terms for timbres and pitch height. We say "sounds like a clarinet though it is too low; perhaps it is a basset-horn". And there are aspects of colour and illumination that only the talented and practiced can master. (In the old days of photography there were a few who didn't need a light metre; I suspect most people are incurably bad at separating saturation from brightness.) So my aim of developing more informative but still vague classifications of pitches might be pushing at a firmly locked door, neurally closed for most of us. I am sure that most psychologists of music, and most musicians who have tried to emulate their colleagues with absolute pitch, would think that this is the case. But I think I have a way of doing it. This essay describes what I think I have stumbled on, late enough in life that my hearing has dulled and my nervous system is no longer plastic. I think that the reason many well-informed people consider the task impossible is that they take the aim to be to emulate perfection rather than a better accommodation to what one can and cannot do. (And, also, I suspect, they are reacting to fraudulent or deluded schemes for developing in many what only a few are capable of.)

For an analogous case with olfaction, see Majid and Burenhult (2014).

2. the components

Begin with the location of sounds. In everyday life this is dominated by the difference in loudness between the two ears, and subtle features of the way one's outer ear reflects sound from different directions, but I am sure these are not the factors that matter for my purposes. Listen to a pure tone, close your eyes so you do not see the source and ask yourself "if it wasn't where it seems most likely to be, what would be my second guess about location?". (Stereo headphones are good for this.) If you were to grab at it in the dark where would you move your hands and where would you be prepared to grab next if there is nothing where you first move? You'll find that the sound has several locations, and the purer the note the more distinct they are. You'll find they move. If you don't hear the multiple locations right away try forcing the sound to move, and note what displacements are easy and which take effort. (Forcing it can be treacherous at a later stage; I suggest this to get you into the wandering pure tones mentality.) You *may* find it is more vivid if you ask for the direction of the sound given the additional information that it is behind rather than in front of you.

Don't invest time in honing this awareness. There are several things going on and it is important to separate them. (I lost months not realizing this fact.) Distinguish two factors, which I shall call *spread* and w*obble*.

spread When you hear a note there is a left-to-right range of where it might be. It seems to occupy all of this horizontal interval. For me B is widest, ear to ear, and E is narrowest. I am curious whether these particular extremes are reached at the same notes for all people. The range is different for different notes, independently of what octave they are in.

Close your eyes and listen through stereo headphones while you hit notes at random on an on-screen keyboard. (It's not hard to make it random: motor control with your eyes closed isn't exact enough to make it anything but, especially if you complicate your mouse movements.) You'll find that the sizes vary between these extremes. Don't worry about making them consistent between trials. That will come, especially given the next factor.

wobble When you hear a note it can move from one location to another, either wandering in space before you or oscillating from side to side. You hear it as being in a certain location and then you swivel your attention to a different direction and you catch it again. Do this with headphones and closed eyes too. For me E rotates through the greatest angle, and B flat through the smallest.

The easiest contrast between spread and wobble is that spread is simultaneous, and occupies the whole angle between its limits. Wobble on the other hand is alternating. It stops in one location before you hear it in another. These differences may be more obvious at high and low pitches, octaves 1, 2 and 6, 7 on the usual numbering. And they are clearer with artificially shaped waves than they are with real sounds. I recommend listening to pitches as sawtooth and as square waveforms (for example using (ii) or (iii) in 12 below), to make wobble and spread, respectively, vivid. To my ear they are undeniable when heard as these waveforms. The waveform labelled 'fifths', presumably because of the balance of its overtones, also gives salience to W, as sometimes for me a synthesised panpipe sound does. It is significant that they remain vivid when one listens with only one ear, for example by wearing stereo headphones and reducing the volume of one ear to zero. Then you can listen to other pitched sounds, piano, flute, and so on, and hear them as sounding like sawtooth or square waves (listening to those components of the sounds) in order to bring W or S to the fore. I find that hearing a sound as an alternation of vowels (AEIOU with a hint of a consonant between them, or the *ojimboway* sequence from 'the lion sleeps tonight'!) brings out wobble, while spread is brought out by hearing it as a single dull indeterminate sound – most unstressed vowels in English – like the sound linguists call a shwa. It may also help to hear spread as if the sound were sung with lips pursed vertically (Edsel grill) and wobble as if they were pursed horizontally (1960s Oldsmobile). (And I find that notes heard with an effort to separate W and S have a special beauty, even when it is not paying off for recognition or memory.)

You may find it confusing when the sound has an inherent oscillation, so that in decaying it seems to be alternately louder and softer. This can sound like wobble. A defence against this is to control the left-right wobble alternation by will, changing its speed or pausing it and the producing a burst of it. (You can do this.) The oscillation of decay is not controllable.

Try this. Play a scale telling yourself that the notes are at first getting 'wider'. Don't define wideness, just listen expecting to find something in the category of size or extent. You should find that the notes get obligingly wider, up to some maximum, and then decline to a minimum, then increase again. The identity of the maximum and minimum should be the same if you go on beyond an octave. Now play the same sequence of notes, telling yourself that they are getting 'thinner'. You'll find that they now get thinner where before they got wider. And the turn-around points, the minimum and maximum, will be near – within approximately a tone – of the previous maximum and minimum. Then try to hold both in your mind at once, best for three notes or so which are uniformly increasing with the one expectation and uniformly decreasing with the other. If you can do this, you'll experience *different* extents increasing and decreasing. So either there are two things going on behind the impressions of width, or the whole business is a matter of suggestion. Evidence against the latter possibility is that you cannot shift the turning points significantly or make arbitrary stretches of increase and decrease, or make different octaves behave differently, or make the impressions differ significantly from one day to another.

In what follows I will describe several ways of hearing a note in terms of spread and wobble, in order to classify, remember, and imagine it. I divide these roughly into the "first method", and "second method", plus some alternatives, though there are a number of separable mental tricks. You may get best results by combining them in some different pattern. An important question that I cannot answer is whether one can skip directly to the method I finally arrive at in section **5,** or whether one has to have taken something like the rambling tortuous route that brought me to it, in order to make it work.

First method. I first listen for wobble, by fixing on a direction for the sound, which can be its apparent direction or straight in front, and then pausing until it disappears there and reappears elsewhere. Eventually this will happen at will when you ask it to, and there will be a jump between the two locations. I do this by stretching out imaginary hands to one location and squeezing the sound with them, until it is confined within a certain area. Then I do the same squeeze for the alternate location. The result is a right-squeeze/pause/left-squeeze pattern. The angle occupied by the squeezes is spread and the angle through which the *centre* of the squeezes jump is wobble. Doing both in one procedure helps prevent either confusing the effect of the other.

One way to learn this may seem like cheating, but it is really just using the results of many comparisons to help with single cases. Play various notes and hear them as having the spread and wobble that they 'should' have. Identify the maxima and minima. If you are like me they will fall into a pattern like this, where blue is spread and red is wobble[2]:

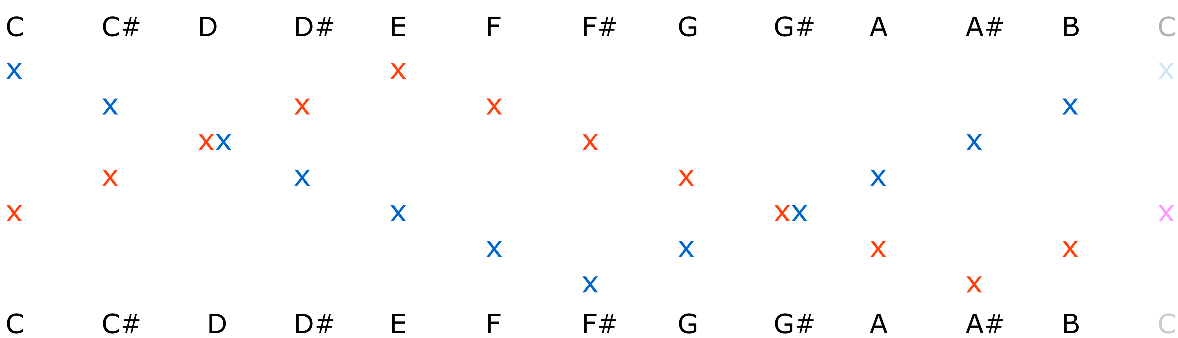

diagram 1

another way of presenting the same information is this.

2 There are pairs of S and W that do not correspond to particular pitches. I suppose it is not inconceivable that these might be evoked by suitable combinations of vibrato and timbre, as 'impossible' colours such as bluish-yellow and reddish-green can be produced by just the right combinations of cone-fatigue and image stabilisation. Then our musical experience would be richer than simple acoustics would suggest. They could also be compared to what are called 'imaginary colours', which are regions of colour spaces that correspond to no possible combination of cone activation. One might speculate that these could be more easily imagined than heard, by imposing the required SW pairs on white noise.

W								
7			E					
6		F		D#				
5	F#				D			
4		G				C#		
3			G#				C	
2				A		B		
1					A#			
	1	2	3	4	5	6	7	S

diagram 2

I have drawn the changes of S and W in diagram 1 S as linear, though I am sure they are in fact more sinusoidal, and diagram 2 should surely be more curvy. But that would be theory over-riding the rough observations I want to present. (But see section **2** just below.) And these are equally tempered notes: to present the data otherwise would be to claim too much for my capacities. The maximum and minimum of each factor are separated by a half-octave, as should be. Not that the max and min of neither corresponds exactly to the minimum and maximum of the other. They are not simple inverses, but the max and min of each is about a tone from the min or max of the other. That is just as well, as otherwise some different pitches would have identical profiles. Diagram 2 makes it clear that no two notes have the same profile.

(The sensitivity graphs of the three colour receptors, while roughly similar shapes, are more out of phase. Taking the range of perceivable colour as an octave – it does roughly double in frequency – the maximum of the medium wavelength receptor is about a third above that of the short wavelength receptor, and that of the long wavelength receptor is about a fourth above the short wavelength one.)

Though no two tones have identical profiles, the close spacing of the two curves does make some discriminations tricky, especially between chromatically related pitches, and between pitches which have similar profiles though the absolute sizes are different, such as D and G for me. Notice, though, as a partial mitigation of these difficulties, that most pairs of notes can be distinguished on generally qualitative grounds: E and G#, for example, have about the same spread but the former's wobble is considerably greater than its spread while that of the latter is less and roughly equal to its spread. Note the subtler contrast between D# and F, though. More quantitative considerations seem to be needed there. (There is evidence that some pairs of notes are harder to tell apart than others, for people with AP.)

Practice, noting spread and wobble until you can tell them apart without consulting the diagram. You may not find you can classify notes quickly, but you should find that you can do so consistently. You will make persistent mistakes, particularly in distinguishing between notes in the two half-way stretches between the extremes, and you may find that you hear occasional extremes of one factor as extremes of the other. Consider how often you classify a note as having a profile near to its correct one, though. You can increase your accuracy with practice. There are suggestions about ways of making fewer mistakes in the following section. (But I suggest that you just skim that section on first reading. You may want to focus your efforts on the second method, in the following section, **4.**)[3]

A warning: I do not know whether the extremes are the same for all people, especially given that I am unsure of my explanations about why this works (see section **11**, 'why it might be so', below.) So it may vary with say, the distance between your ears and how strongly one ear, eye or hemisphere is dominant[4]. After

3 I shall say little about key perception in this essay. I think the slower first method works better for identifying keys. You scan left to right, keeping your attention directed at the music, and noting the size of the jumps between moments when you seem directed right at it. Then within each such location you squeeze a spread that confines the music to that location, and which itself jumps to the next. The first gives the wobble of the tonic and the second its spread. Bear in mind that the spread may be wider than the wobble, and thus extend to the next location. (There may be a horizontal series of locations rather than an alternation between two.)

4 This is a very basic question. If there are individual differences, it is tempting to look for correlations with those revealed by Deutsch's tritone illusion, which shows systematic individual differences in which of two tones of indeterminate height (Shepard tones), separated by half an octave, is heard as higher (Deutsch 2013b, pp 142-144.) The connection would be via the testable suggestion that people use a heuristic of relying on size of interval to determine which of two notes within a half octave is higher, and when they are about equally spaced choose either in terms of S or in terms of W. There would be two determinants of individual differences: where the extrema are and which factor was used. The fine structure of judgements, including where uncertainty is found, would distinguish the use of the different possible heuristics.

all some colours look green to some people and blue to others. But it should not be hard to adapt my picture to the way things sound to you. You may find it helps to construct your own diagram analogous to the one above, putting together a number of pair-wise comparisons of neighbouring notes, using sawtooth and square wave generators if that helps, so that the pattern has one maximum and one minimum for each of spread and wobble separated by a half-octave, with the maximum of each being within a tone or so of the minimum of the other.

2b. A more precise diagram

Sawtooth and square wave-forms, and no doubt further experience, have given me more confidence in my judgements of S and W than I had at first. I think the deviations from linearity are fairly small, and confined to the upper end of the curves, because we find it hard to perceive a wobble of a full 180 degrees, and the bottom end, because we would not notice a spread of 0 degrees. (But then sinusoidal patterns will always diverge from linearity more at their extremes.) I have also dissuaded myself from my original sense that the extremes were somewhat sharp of the named notes, though it seems very unlikely that they coincide exactly with concert pitch names. Below is a modification of diagram 1, in which both S and W are measured in terms of angles. W presents itself naturally as angles; S less so. Instead of drawing curves – and leaving more hostages to fact – I have left the lines straight and modified the scale at the left, indicating in square brackets the values that would follow from full linearity.

To think of S spans as angles imagine your hands held out at arms' length, hold the span with them, and see how much of a clock face that makes. The Greek letters in the 'corners' column are the measurements of the four extreme values. (Diagram 2 shows why I call them the corners.) They return in section **7 (ii)**.

Note	C	C#	D	D#	E	F	F#	G	G#	A	A#	B	
[180] 170-	x				x								α
[150] 135		x		x		x						x	
120			xx				x				x		β
90		x		x				x		x			
60	x				x				xx				γ
[15] 30						x		x		x		x	
[0] 10							x				x		δ
degrees													**corners**

diagram 3

3. avoiding mistakes

You will find that you sometimes confuse the extremes of the two kinds, judging wide spread for large wobble, and vice versa. And sometimes you will get the absolute sizes wrong. Here are three ways of minimizing mistakes.

Wait for the pause Wobble alternates and spread runs through the positions between its end points. So wait for the pause between wobble locations and make sure that each spread stretch occupies all the space between its end points.

Find the middle A wide wobble can stretch one's judgement of spread and a wide spread can stretch one's judgement of wobble, because one is inconsistently assessing the angle between extreme right and left or between the two inside end-points. The solution is consistently to assess wobble from the mid-points of the squeeze and spread from its end points.

Don't pre-judge it Especially if you are doing too many in a series, you can form judgements based half-consciously on what one takes the note to be. This can result in interval judgements infiltrating pitch judgements. So make practice sessions short, especially at first. (Four is good; but it is hard to stop at four, especially when you are making mistakes and your pride requires that you get enough right.) This third factor combines with a fourth: we get tired. The thinking and imagery here is new, and one is learning something of a kind one has never learned before, so only so much is possible before one starts to deceive oneself about whether one is really following the instructions.

4. Pitch qualities

Diagram 1 has more information than one is likely to keep in mind when trying to identify a note. In this section I show how we can get classifications of pitches from it, which are more similar to colour words. I present them as illustrations of what can be done, not as representations of concepts that we actually apply to sounds in music. (For colours we know that some classification systems are psychologically more natural than others, easier to understand and use. See Berlin and Kay 1969. I expect something similar is true for pitches.)

One division into vague pitch qualities, that is linked in an obvious way to diagram 2, can be illustrated as follows:

top W, top S	top W, high S **I** **E**	top W, low S	top W, bottom S
high W, top S	high W, high S **II** **D#** **D**	high W, low S **III** **F**	high W, bottom S **IV** **F#**
low W, top S **V** **C**	low W, high S **VI** **C#** **B**	low W, low S **VII** **G** **G#** **A**	low W, bottom S
bottom W, top S	bottom W, high S **VIII** **A#**	bottom W, low S	bottom W, bottom S

categories A

Another categorization is

+ S greater than W "saturated" C, C#, A, A#, B	= same S, W D, G#	- W greater than S "unsaturated" D#, E, F, F#, G

categories B

Combining categories A and B, we find

C:	V+	F#:	IV-
C#:	VI+	G:	VII-
D:	II=	G#:	VII=
D#:	II-	A:	VII+
E:	I-	A#:	VIII+
F:	III-	B:	VI+

In this way each note gets a unique characterization. It seems to me very likely that there are simpler pairs of categorizations that also deliver uniqueness. Most important would be categories that fit the way notes sound once one has learned to make these discriminations.

All the categories in A are vague, except those denoting maxima and minima, and one can expect there to be borderlines of uncertainty between them. The same goes for the + and – categories in B: tones with greater differences from the extremes of S and W will be easier to classify. Maxima, minima, and the "=" category in B, though in theory precise will also be harder to apply to cases near their borders. Practice in applying both categories is, I think, useful mostly as a way of training the distinction between S and W. It also helps cure you of the idea that if you cannot name a note by letter then you are not perceiving its character as a pitch.

A contrast between the methods I am describing and some methods purporting to develop absolute pitch is worth noting. These methods often tell one to listen to a single note and concentrate on it until one can recognize it, and only then move on to another. So all going well one gets single notes one at a time with full accuracy. I have doubts about how well or often this will work, and indeed about the honesty of many such enterprises, but I have not seen quantitative studies of them. As I remarked, my emphasis is on access, letting accuracy come as it will. But it is also global: the conscious judgements of the two qualities come in degrees, and can be refined with respect to a number of notes simultaneously. In fact, it is best done for several at once, as the comparisons of spread and wobble often give information about the character of a note. ("It has quite a wide spread, but not as wide as E, so it is D or F.") Doing it this way allows memory for intervals to scaffold memory for pitches. This is possible because we can use an approximate grasp of any of three factors – W, S, or interval from a previously identified note – to give an approximate grasp of the others, which can then refine the original estimate. On the rival approach there are no descriptions suitable for approximation, nothing that comes in indefinitely small increments, to refine and set up a mutual nudging[5].

5. doing it quickly: second method

I have described listening to tones until their spread and wobble emerges, using various tricks to make them emerge and keep them separate. But this is slow. It works less well with brief notes or notes with a rapid decay. And you may need to be convinced that we are dealing with the processing of pitches themselves rather than an acoustical phenomenon. So it is important to speed up the process by holding the note in your short term memory and dealing with it there. The following procedure works for me. It allows quicker judgements, and I can refine the sound

5 This might explain why some people with absolute pitch find music that is slightly out of tune or not at concert pitch more disturbing than non-AP people with good relative pitch do (Deutsch 2013). The explanation I suggest would trace this to learning a range of discrete notes and not having a good grasp of the continuous variation beneath them. People with a more global grasp of pitch, on the other hand, will be able to hear a note as a certain perhaps approximate degree of spread and wobble, which allows it to be grasped even if it is not one of the perfect exemplars. Of course people with really wonderfully accurate AP will be able to do this too ("that's how A would have sounded in 1860") and there are stories of such people rejoicing in exotic mis-tunings.

of the pitches to enable me to hear the note *as* having the relevant attributes.

Hear notes as beginning at a definite point in time but having two endings. The note begins with a percussive or consonantal onset sound, which soon ends. Hear the end of this onset as having a sound quality itself. (Remember the sudden palpable silence when a noise that you have become unaware of suddenly stops. Think of the mouth-closing sounds that English speakers put at the end of French final vowels.) You will find that the on-sound and the off-sound occupy different locations. (For me it is a left to right skip.) That is wobble: the separation of beginning of *on* and end of the first *off*. But simultaneous with this, or beginning a milisecond later[6], there is also a continuous spread, not a skip, to fill an extent that can end either within or beyond the wobble extent. That is spread.

So: hear notes as having a beginning, an extended percussive/consonantal onset, and two endings. Practice with known notes – the 'four corners' exercise in section **7 (iii)** below is good for this – until you can hear these, and then practice with unknown notes until you can at least approximately classify them. As you do this you can begin to chip away at two further, linked, projects. (a) getting the representation of a note to summon an aural image of the note; mind's eye prompts mind's ear. (b) getting representation and sound to fit, so that you can hear a note as having its stretch and wobble. It sounds like this combination, much as a specific shade can come to look like a combination of a particular saturation and a particular brightness, and then that combination presents itself as the way the colour has always looked. I think that paying attention to these further tasks helps with the seemingly less demanding recognition task, as learning the exact attributes of given notes is a good way of fixing them in your mind so you can recognize them. Short term aural memory combines with long term visual memory to foster long term aural memory.

---

6 It would make sense that S should appear later than W, on the gloss I gave in introducing them above. If W corresponds to finding alternative directions to search for a sound source, in each of which one then carries out a search of a given spread, then one will often fix W before S.

As you do this you may find that the first stretch – W – and the second – S – begin to sound different, and produce different visual images. The first is for me often darker and more sour while the second is lighter. The first sometimes sounds as if a dark vertical line were moving rapidly across the extent, while the second usually sounds as if it were spreading out to fill its space. The first sounds as if it ends abruptly when it has skipped to its end while the second ends in a sower more rounded way. These are likely to vary from person to person, so attend to the forms they take in your experience.

6. alternatives I arrived at the second method after trying several refinements of the first method. This presents a decision to the reader. The second method is easier and better than the first, for me. Should you jump straight to it, or does it work only after you have used the series of refinements I made between the first and the second? You'll have to find out for yourself. In this section I do not run through the refinements but discuss three aspects of them. They may help you find methods that work for you, or bridge the gap between my slow but straightforward first method and my better but more mysterious second method.

touch I combined spread and wobble into a single visual/tactile image. Both are horizontal, but S is a continuous extent that you can mentally squeeze with imagined outstretched hands. (Fixing on a standard distance away, most naturally at arms' length, helps resolve issues of scale.) W is an alternation between two end points, superimposed on the S extent. I think of it both as a glancing from one side to another and as a little ball disappearing at one side and appearing at the other. The crucial thing is that you do not swing the S extent from one side to the other, but hold it before you with the W oscillation occupying the same horizontal line but differing in length, longer for some notes and shorter for others. You feel as much as visualize these, but then remember them as spatial distances and angles. Practice with extreme cases, particularly with square and sawtooth waveforms, will provide a basis from which you can extend to more subtle ones. (Aristotle thought that touch was the basic sense – *De Anima* II, 2 – and whatever we make of that thought, touch does connect perception and action fairly intimately, since you cannot separate feeling objects from pressing against them. And unconscious perception can be correlated with actions that we can then be conscious of.)

three dimensions In struggling with aural memory and classification I would frequently find three dimensional visual/tactile imagery suggesting itself. I would take these at first as important indices of perceptual properties of pitches, and then I would find that vivid as they were they did not help much with the tasks I had set myself. For all that, forming these images may have had some effect on me, and the impression persists that this unreliable phenomenon goes to the heart of things. So I will try to describe them.

What I am describing as a horizontal W alternation sometimes presents as a separation along the axis in front of me, as if W extents of different sizes were different distances away. It can also seem that the larger spreads, for both W and S are nearer and the smaller ones further away. Both of these would make sense if they turn on registering phase differences between ears or, in one channel, between samples taken at successive times. They would then be parallel to the use of binocular disparity in depth perception.

scale Some notes sound very small. A faint and distant sound at first presents itself as having next to no extent, so nothing to extract S and W from. But there are really only three scales: the punctuate perceived location, a one-sided hemisphere, and a two-sided panorama. Decide on one of the latter two, preferably the last, which is easiest if it is centred straight ahead, and insist on it. Then the larger extents reveal themselves.

I expect there are many variations on these, and that some have more effect for some people. The ingredients are separating the two components and making use of both three dimensional imagery and motor/tactile imagination. Any combination of these is worth trying**.**

7. a series of exercises

(i) *Preliminary.* Take three notes. (I worked with E, A, C, because of a website with these three as digital tuning forks. See (i) below. But they are contrasting pitches on my classification.) Play one, hear it in terms of S and W, and combine them until they merge into the way it sounds as a musical note. So you are now holding the sound in your mind but hearing it as consisting of the two components. As it begins to fade in your aural memory use the two components as ways of bringing it back. When you can do this reliably, summoning it when it is no longer in consciousness, do the same for another, then the third. Lie in bed willing them into mind, in various orders to make basic three note tunes. Then when this is solid – the next day, perhaps – go back to the source of the notes, first summon them mentally as vividly as you can, and then see if they fit the notes when you play them.

(ii) *The four corners*. The notes with extreme values of W and S are for me C (max S), E (max W), F# (min S), and A# (min W). (If they are different for you, modify this exercise.) The other values of the profiles of these notes are intermediate between the extremes. Thinking of both S and W in terms of angles extended or oscillated through, we have four values. Call them alpha – about 170 degrees or just less than 9 to 3 on a clock-face; beta – 90 degrees or 12 to 3; gamma – 60 degrees or 12 to 2; and delta – 10 degrees or 12 to less than 1. See the remarks in section **2b** on measuring S in degrees. As follows:

	C	E	F#	A#
W	Gamma 60	Alpha 170	Beta 120	Delta 10
S	Alpha 170	Gamma 60	Delta 10	Beta 120

It helps to learn the angles as felt rather than as calculated, so that the correlations with these pitches are right. (Especially for alpha and delta where the suggested departures from linearity are surely not precise.) And it is important to remember the notes in terms of the sizes of the angles rather than as remembered tones. For most of us the latter will not work, at any rate until pitch sense and awareness of spread have become thoroughly integrated. Try spreading your hands and arms to give a tactile grasp of the angles. They can then be used as guidelines for the spreads of other notes between them in size. Note that before the corrections for non-linearity at the extremes (see diagram 3 above and the discussion in **2b**) these descend in 60 degree steps: before correction 180-120-60-0.

When working on this exercise and singing the notes (see section **9**) I found myself feeling the angles as spreads within my mouth, with alpha being the full mouth width at a level with the line between my ears, and the others correspondingly folded angles from the centre-point of that line. (So now three dimensional imagery is returning.) I found that this gave me a good memorable hold on them[7].

Produce these notes and fix the sizes in your mind. (Together, they sound rather Debussy-like.) Again, sawtooth helps with W, and then the angles can be remembered and applied to S. Choose two – say C and E – and hear them as combinations of these precise values of S and W. Do this for 24 hours until you can summon these two notes at will. Be patient. Lying in bed and sitting in the sun, letting it take its time. When you are confident of these add the other two. (I find E the easiest to summon. Then C is E backwards, with W and S reversed. F# straddles the S component of E, and A# is F# backwards. Don't add them too rapidly: each additional note will take at least two days.) Then summon them in arbitrary orders, as little unexciting tunes, E, C, B#, F#; F#, C, B#, E, and so on. Don't move on till you can do this.

---

7 Notice how three-dimensionality returns mediated by imagined touch. I find that I can summon the notes by rehearsing the sequence of imagined motions. In making the notes emerge from the spatial gestures I have to be alert both to inhibiting my diatonic instincts and to accommodating departures from linearity. In particular, if I do the sequence in ascending order I have to push down on both components of F# making them slightly more distant, and lessen the distance between them, to prevent F# from sounding too close to F.

A crucial moment in this process comes when the notes that came easily not long ago hide. Perhaps they were vivid and effortless when you went to bed and in the morning when you request their presence they refuse to appear. Don't rush to keyboard, computer, or tuning forks to find them. Instead try this. First remember the desired S and W as purely spatial. Then think of a pitchless sound, the buzz in your head or an imagined hum, or an imagined voice or sound some way distant, and combine it with the desired S and W. Impose the S on it and hear it, having determined in advance the W that is going to apply. Make the sound take that spread, keeping the wobble in the back of your awareness. You will find that the sound assumes a pitch. *Now* go and check whether it is the pitch that you wanted. When you can do this reliably you know you are making progress.

(iii) Now add D and G# to complete the whole tone scale that the four corners are most of. These two notes have in common that their W and S are the same, but for D it is at Beta and for G# it is at Gamma. The extents or angles are as in diagram 4 below.

C alphagamma	D betabeta	E gammaalpha	F# deltabeta	G# gammagamma	A# betadelta
α β γ δ δ γ β α	α β γ δ δ γ β α	α β γ δ δ γ β α	α β γ δ δ γ β α	α β γ δ δ γ β α	α β γ δ δ γ β α
	both at beta	opposite of C	both shrink by 1	both at gamma	opposite of F#

diagram 4

(iv) Now you have to learn the rest. But that's not as daunting as it seems, since (ii) and (iii) will have fixed in your mind the four 60 degree spreads, so you can now interpolate the 30 degree ones, remembering which ones, in which S-W combinations, correspond to which notes. Not trivial, but manageable.

(v) *an auditory-tactile-visual experience* When I was working on (iii) I made a preliminary version of diagram 4, to fix in my mind the sequence of separations in the whole tone scale that grows from the four corners. I rehearsed using the diagram to summon the tones, as described, for an evening and went to bed. Then for two hours I felt as if I was seeing and touching the notes, as if they were vividly and tangibly before me as large spatial shapes moving around me. The shapes were silent, but attending them made it easy to summon the notes in my imagination. In spite of the vividness of the shapes I could not determine which were their relevant geometrical features. This continued in hallucinatory fashion till I finally slept.

8. Imposing the patterns

Because this structure gives a single image rather than a pair, it can be superimposed on a variety of sounds. In particular, it combines nicely with white noise. Do this: turn on a white noise generator. (There are several on the internet: II find the variant called ‘brown noise’ works best for this purpose. But an electric fan will do.) With that in the background produce some clear pitches. Fix in your mind the spread and wobble of the pitches. Then superimpose this shape on the white noise: hear the attributes onto the noise to make it take the pitch. Hold the pitched noise after the clear pitch has faded, and then bring it back again. Do this with sequences of pitches. (Best if they do not have a structure of easy intervals, to prevent yourself cheating, as for example in section **7 (ii)** and **(iii)** above.) What I find is that doing this produces vivid echoes of the original pitched notes, with their original timbre, and when this fades the pitches can still be recovered from white noise considerably later by thinking of the intended structure as a way of summoning the exact pitch-shape. The aim is to expand the number of notes that you can summon, until you have a mental piano you can play at will[8].

The experience of hearing these ghost pitches emerging from white noise is quite different from that of remembering a melody. It is more ‘external’, less memory-like

8 And then comparing heard notes to its output, classifying them by 'analysis by synthesis' (Bever & Poeppel 2010.)

(as if there is a phenomenology of the reconstructive aspect of memory) and more neutral as to timbre. You can superimpose it on un-pitched timbres, in particular on the high white noise of blood circulating in one's head that many people can hear. (I nearly always can, and I suspect everyone can if the background noise is subdued enough) So one way of summoning a pitch at will is to impose the relevant shape on this home-made white noise. (I would not recommend concentrating on summoning any single note in this way, rather than the skill of summoning one from a variety, for fear of inducing tinnitus. You might turn into Robert Schumann[9].)

You can progress from loud white noise, to softer white noise, to pitch-neutral sounds around you. And eventually to the gray screen of your imagination. It takes a lot of repetition. An important step occurs when you manage to hear the two attributes as part of the phenomenal quality of notes. This takes a certain amount of persistence and experimentation. I think you have to stumble for yourself on whatever mental trick is needed. (For me, the trick has something to do with thinking of wobble first and spread second, and hearing the transition between the two as producing what the note sounds like.)

I find that although I am good at distinguishing musical instruments by timbre when I hear them, and as an ex-oboist (never a very good one) I can sometimes distinguish national styles and schools of oboe tone, I find it considerably more difficult to imagine particular instruments with any vividness. For me, various un-pitched noises are easier to imagine. A list of typical noises is:

a wooden ball dropping onto a wooden floor
a small metal object dropping onto a hard floor, such as a stone one
a pebble dropping onto a stone
water dripping into water
impact of metal on hollow metal poles, e.g. scaffolding poles

9 Remarkably many composers have developed tinnitus. A quick search produces Beethoven, Janacek, Boyce, Holzbauer, Draessiker, Fauré, Smetana. I conjecture that a common factor is excessive attention to the sounds in ones head.

a knock on a wooden door
footsteps on gravel
some unbeautiful bird cries, such as seagulls and crows
-- and nearer to determinate pitches --
wind noises, including the sound of blowing over the top of a bottle
a wineglass being rubbed around its rim

I expect everyone can produce their own list. I can imagine a series of pitches from white noise, and then superimpose them on sounds like these, to produce 'instrumental' versions of the sounds. One can see how with a lot of practice this could develop into an orchestrator's imagination.

9. Singing It is easy to deceive yourself that the note you have just imagined is the note you are now playing, rather than your imagination having shifted to match the sound. The obvious defence is to sing the note you imagine and compare it with the note you hear. I have never been good at singing remembered and imagined notes. (Strange and disconcerting moments when I ask someone to remind me of the identity of some famous tune but they do not recognize it from my attempts.) So learning to let the S and W characterization of a note guide its production is not an idle extra.

I find that thinking of the two component description while reproducing a note helps. Since beginning to do it I am rather less hopeless. The core, I'm sure, is just to keep the two separate and in mind while producing the note, trying to match them and not worrying about making a nice sound or doing anything closely resembling singing as one knows it from good singers. (Perhaps the aim of singing like these people gets in the way. In fact the sounds that I find most easy to match notes with seem to me like ethnomusicological examples of vocalisation from cultures not focused on our ideas of tuneful singing.) Then eventually one will train oneself to do it. There must be many distinct elements of vocal control involved, and one can't expect to be aware of many of them.

For me, it helps to keep in mind the double-ending representations of the second method (section **5**) and to sense them within the resonating areas of my mouth. No doubt the effect is an illusion, but it allows me to direct my fine-tuning to the two separately. I find this works best if the noise I produce is a whisper or a hum, nothing like singing to an audience[10].

10. other sounds There are combinations of S and W that do not correspond to notes. (See footnote 2.) A waste? Perhaps not. Not all sounds are pure tones, so we can ask whether some of these un-named combinations correspond to sounds that are not musical pitches. I am very unsure of what I say on this topic so this section is best taken simply as a description of a way that one can hear. I

The same sense of alternation and extent can be had from many non-musical sounds as from tones. It is harder to pin down what the exact values are, no doubt because most natural sounds are so complex. The simplest complex sounds are chords, combinations of pure tones. (Of course we can make chords of even a few notes that are pretty complex as sounds, by combining rich timbres with many overtones.) So begin with them.

Standard intervals such as 3ds, 4ths, and 5ths do not have uniform simple characterisations in terms of spread and wobble. Semitones do: a semitone is one step away(30 degrees, in the terms I used in **2b**) in both S and W. And if two notes are separated by exactly this amount in both components then they are a semitone apart. Octaves obviously have the same profiles. But both fourths and fifths can have (W, S) patterns of (2n30, 2m30) for odd n,m, n,m > 1, n≠m, and tritones

10 There *might* be a connection with the well-known phenomenon that people often sing well-known songs in the keys of their famous recordings, though they could not name the keys. The connection would be that people have motor representations of what it would take to reproduce the song.

(half-octaves, diminished fifths) can have patterns of (n30, m30) for even n,m>1, n≠m. It is not surprising that 4ths and 5ths share a profile, as when you ignore tone height they always come as a pair. (If you play Shepherd tones of say C and G they will be related as both fifths and fourths.) If chords are played simultaneously, rather than as arpeggios, the two components of each note are audible, so that four or more extents can be heard. When they are minimally spaced, as with semitones or very widely spaced, as with most tritones, the effect is confusing and hard to process. One aspect is likely to be that the implicit spatial information is not clearly differentiated enough.

Consonance is a many-sided affair. There is the collision of overtones; there is the susceptibility to resolution in a tonal context. And there is the digestibility of the implicit spatial information. They are all different, and it is the last that concerns us here. Two-note chords whose S and W are can be divided into stretches of between three and five times 30 degrees seem to be easy to hear, perhaps because when the S or W of one note is near to that of another they can fuse while that of the other is further away, thus preventing ambiguity. Pairs both of whose components are close are harder to hear. Many fourths and fifths contain a 30 degree separation between either the W or the S components. (Never both.) When you hear them, listening to the four spreads – and thus not separating the chord into its component notes – the effect is of the two near S or W fusing and of a rolling sound between the two more widely separated ones. You can concentrate on the fusion and the rolling, so that they are sounds in their own right. They are not tones with names, but they are sounds that are familiar from music and everyday life.

In this connection it is worth mentioning that when you pay attention to these aspects of combinations of notes, not all chords spanning the same intervals sound the same. For example the two tritones D-G# and G-D# sound different in that the former tends to fusion in W and roll in S and the latter to roll in both. Similarly of the tritones C-G# and E-A# the former is more fifth-like in that there is more fusion in W and the latter is more fourth-like in that there is more fusion in S. Similarly the two fifths C-G and E-B# sound different in that there is more W fusion in the former. The latter sounds like a more complex chord. (I can imagine mistakenly thinking there were more than two notes involved.)

Triads on many notes have a curious lining up quality. The W and S points of the components coincide, sometimes W with S and sometimes W with W. That would make it easier to base key recognition on note recognition, and would explain the difficulty of telling complex sounds from pure sounds, especially for non-musical but pitched sounds such as telephone beeps. They sound like notes, but when you try to decide which ones you cannot. A C major triad, for example, sounds to me as if it has the S of C and the W of G. I can imagine that if it was composed of sounds with many overtones I might think it was a pure note with a thick timbre, before finding I could not identify it. I wish I understood this better. I suspect that it requires a more accurate version of diagram 3. (But I also suspect that it is something at which the musically trained have a peculiar disadvantage: they have learned not to hear things this way.)

I expect that some will find these descriptions of these sounds more apt than others. Beneath the uniformity produced by musical training and practice there must be a greater variety in what comes easily to different people than we can express in terms of success or failure at tasks presented by our musical tradition. Perhaps considerations like those of this section can help show how many varieties of ears we jointly possess.

All my life I have felt that some scenes look like sounds. These are usually scenes with some periodic scattered aspect. Moving highlights on water, light filtering through a canopy of leaves, moving shadows, vertical blinds moving in the wind. (There is a Japanese word *komorebi*, meaning 'The scattered, dappled light effect that happens when sunlight shines in through trees'.) Since beginning to hear spread and wobble, more sights prompt a sense of imagined sound, mostly objects moving against a background (people in a park), bright three dimensional arrays in a dark space (arrays of stars seen so that some stand out as nearer and some as further away), and gaps between approaching and receding objects. (Driving towards trees, seeing their branches as separated in receding distance and these separated gaps as moving towards you.) All of these can be seen as evoking my two qualities, as moving spatial processes. (I think motion is essential. Hearing and sight are related in different ways to time. Everyone knows this and it is part of the difference between music and painting, or even music and silent film. But it is hard to state in a way that is more than hand-waving.) But to my initial surprise not many of them gave any sort of representation of pitches. I now suspect this is because pitch is only the surface of a much larger phenomenon.

11 why it might be so

Why should these qualities of sounds be correlated with these visual-tactile qualities? Here are three possible explanations, one philosophical/psychological, one biological/acoustic, and one neurological. They are not incompatible, and they may all be wrong.

The p/p explanation is simply that there is a systematic structure of contrasts between pitches, and so there is among images in motion, and if we fasten on contrasts which are structurally parallel then any representation of the one gives the possibility of transfer to the other (Rosenthal, forthcoming). But why images in motion? Why S and W, which do not correspond to any musically establshed features of pitches? If this explanation is to succeed then there must be more to say.

The b/a explanation builds on the fact that a pure note of a given pitch will interfere constructively with itself at one ear and destructively at the other, if it is at a series of locations with respect to the hearer. (It depends on the pitch and the separation of the ears.) So information about which directions of one's head make a sound take its maximum generates information about possible locations of the source. For pure notes there is a multiplicity of locations, but this is much reduced for a complex sound, since there will only be a few locations which occur in the possible locations of all its pure components. It is thought that birds, especially owls, locate sound sources in this way (Saberi and others 1999, Wagner and Frost 1993.)

(It is suggestive that the evidence for phase-based location, which works best for complex sounds, is mostly from birds. (Saberi and others 1999, Wagner and Frost 1993.) Small vulnerable birds, that want their presence but not their exact location noticed, tend to call in pure notes, while larger and tougher birds tend to have much rougher calls. No more than suggestive, though.)

If a mechanism along these lines is found in human hearing, then the connection between pitch and spatial location would be much less arbitrary. But this particular connection? I am not going to work this out in more detail. It would be speculative, and it would distract from the point that any such explanation could be wrong while the technique of conceptualizing pitches in terms of width and wobble works, allows us to make discriminations and educate our imaginations in ways that we could not previously.

The suggested mechanism also suggests some possible neural correlates between individuals with and without accurate pitch perception. The contrast between S and W is most vivid when one hears width as a simultaneous squeeze from both sides and wobble as an alternation from side to side. (That would tie to the spatial depth phenomenology too.) People with the ability to alternate attention rapidly from one side to the other – corpus callosum? – would be better at this.

The usefulness of white noise for learning pitches has a very plausible outline explanation. Given a scene in which there are many simultaneous pitches and timbres – a movie jungle soundtrack – a person or other animal will want to be able to locate a particular source in the array. The scene provides the white noise and a briefly noticed or remembered sound provides the source that one needs to separate from the background. So one needs a way of operating on an undifferentiated jumble of sounds to make candidate locations for a target sound salient. Then one can look or reach to that location. And if one misses one needs to know where one's fall-back location is.

There is a third rather different possible explanation. The comparison of periodic information through different channels might be a fundamental neural process, occurring in many different contexts. Then there would be similar neural processing of information in different modalities, which could allow one to be compared profitably to another, if analogies at the right levels are created. (I find this one attractive. It is not incompatible with the others. But it would be hard to find direct evidence for it at the present state of our knowledge.)

12. appendix: actual and possible technology

To practice these things you need to be able to produce named notes and listen carefully to them. You will need to test yourself see how near to getting them right you are. I have found these useful:

i) Online tuning fork: http://www.onlinetuningfork.com/ E, A, C

ii) "Absolute Pitch" app for android. (Don't be put off by the title; that's just to make sales.) I have produced random notes by closing my eyes and sliding the mouse up and down the keyboard. Imperfect motor control and friction conspire to make the eventual note unpredictable. I like the variety of timbres available.

ii) Virtual keyboard:
http://www.bgfl.org/bgfl/custom/resources_ftp/client_ftp/ks2/music/piano/
There are more grand piano-like internet keyboards, but I like the variety of timbres available on this, and the chord mode.

iii) an online waveform generator:
http://www.wavtones.com/functiongenerator.php

You can do it all with these. Or others. But there is room for software targeted specifically at the method I am describing. A simple useful program would present choices of notes to distinguish, increasingly many as one had mastered smaller sets. For example one could progress from the four corners of **7(ii)** to the corresponding whole tone scale, and eventually to all twelve. This would be best with tones of indeterminate height (Shepherd tones) to prevent one learning the notes melodically.

I would be interested to discover the uses of what one might call a well-tempered white noise machine: an ever-changing combination of notes from across the concert pitch chromatic range. (Like banging an unevenly shaped stick rapidly and repeatedly against the strings of a piano: you won't get anyone to volunteer their Steinway for this. I had to make do with the chord mode of the virtual keyboard mentioned above.) This would be useful for training oneself to hear pitches on demand. One would hold the specifications of a note in mind and then let it appear from the noise. This can be done with regular white noise, but then there are many more unwanted targets that one has to hear past.

There would be uses for wave-form generators that do what I have used sawtooth and square waves for, to make W and S evident. There might be wave-forms that made each of these even more evident. There might be wave-forms that made both simultaneously evident and different. Then one would progress by little increments to more musical-sounding notes. This would combine naturally with the 'equally tempered white noise machine' described above.

A radical idea would be a sound generator that presented a particular spread or wobble. It would give a white noise-like array of pitches – best along the lines of the well-tempered machine described above – which all had in common a particular S or W. Then one could hear these in isolation from hearing a pitch. They would be like Shepherd tones (Deutsch 2013c), but more specialised. One could test oneself: what is the S or W here? And one could demonstrate the components: S is what all these have in common, and W is what all these others do.

Bibliography

Ali, M, Klyne, M (1985). *Vision in Vertebrates*. New York: Plenum Press

Bever, T and D Poeppel (2010) 'Analysis by Synthesis', *Biolinguistics* 4, 174-200

Dennett, D.C. (1991) "Heterophenomenology" in *Consciousness Explained*, Penguin Press.

Deutsch, Diana, ed. (2013) *The Psychology of Music, 3d ed.* Academic Press.

Deutsch, Diana (2013b) 'Absolute Pitch'. in Deutsch 2013, 141-182

Deutsch, Diana (2013c) 'Grouping mechanisms in music'. in Deutsch 2013, '183-248

Gerstner and others (1996) A neuronal learning rule for sub-millisecond temporal coding *Nature letters 383,* 76-78

Hofman and others "Relearning sound localization with new ears." *Nature Neuroscience* 1.5 (1998): 417-421.

Kristan, William B. (1998) "He's got rhythm: single neurons signal timing on a scale of seconds." *Nature Neuroscience* 1.8 (1998).

Majid & Burenhult (2014) Odors are expressible in language, as long as you have the right language. *Cognition*

Rosenthal, David M (forthcoming) 'Quality Spaces and Sensory Modalities'. In *The Nature of Phenomenal Qualities: Sense, Perception, and Consciousness*, ed. Paul Coates and Sam Coleman, Oxford: Oxford University Press.

Ross, Alex (2007) *The Rest is Noise* 1st edition , Farrar Straus and Giroux (the pagination seems to be different in other editions, the quotation at the head of this essay is on p 168 of this edition. Ross gives M Huttunen in *The Cambridge Companion to Sibelius* as his source, which I haven't checked.)

Saberi and others (1999) Neural bases of an auditory illusion and its elimination in owls *Nature Neuroscience* 2, 656 - 659

Wagner and Frost (1993) Disparity-sensitive cells in the owl have a characteristic disparity *Nature* 364, 796 - 798

Websites on stereograms:

http://www.eyetricks.com/3dstereo.htm

http://en.wikipedia.org/wiki/Autostereogram