Nasals are the colour of the morning sky: a pilot study of the synaesthetic correlates of distinctive features in phoneme-colour synaesthetes

Candidate number: 9553F

March 13, 2017

Word Count: 9811

Contents

Abstract

This study explores the possibility of using phonemic synaesthesia as a source of evidence in phonology: if synaesthesia is sensitive to phonological units, it can serve as a diagnostic tool for phonological phenomena. Synaesthesia has been shown to be able to refer to acoustic properties and morphological units, which entails it has access to linguistic information. However, there is no similar research on the interaction of synaesthesia and phonology. The literature mentions the existence of phonemic (as opposed to graphemic) synaesthetes, which suggests that synaesthesia can make reference to phonemes: if /f/ is blue, that will be the colour for the coda segment of LAUGH lɑːf, even though <f> isn't in the spelling. It is therefore not unreasonable to suggest that it may also be influenced by smaller phonological units such as distinctive features, which might trigger colour percepts corresponding to feature-based natural classes in phonemic synaesthetes.

The study has mapped the synaesthetic colour correlates of 3 phonemic synaesthetes for items in their phonological systems and concludes that phonological synaesthesia seems to be influenced by phonological features. These results are of relevance for certain larger debates in phonology, such as whether features are innate or emergent, whether they are devoid of phonetic content or not, whether features or rather gestures are the building blocks of phonemes.

1 Intro

Research on the interaction of synaesthesia and linguistics is scarce, and what does exist focuses overwhelmingly on graphemes. There appear to be no detailed case studies of phonemic synaesthesia, an area which seems to be an accidental gap rather than a systematic one. Nonetheless, individuals with phonemic synaesthesia reportedly exist, and because their synaesthesia is based on psychologically real units of phonology it is a potent testing ground for theories, as colour correlates are easily quantifiable and as such can reveal a great deal about issues such as category membership and underlying representations, which would be represented visually in the form of similar hues or luminosity, and can shed light on the larger linguistic debates. For example, proponents of Distinctive Feature Theory would expect all phonological effects to be categorical and all phonetic ones to be gradient, while in a purely connectionist approach all effects should be of the gradient type. Each theory makes different predictions for what the data of a phoneme synaesthete might look like: connectionism predicts all tokens in a phonetic space to change colour gradually when moving across one dimension, regardless of phonemic membership. DFT on the other hand predicts that not only will the colour changes be drastic from one phoneme to another, but also that the colour patterns which appear can be modeled on phonological features.

The aim of this dissertation is to conduct and report on a pilot study investigating whether distinctive features are psychologically real, using as a measure whether or not they influence colour patterns in phoneme-colour synaesthetes, and how phonemic synaesthesia may be used as a diagnostic tool in phonology. The study finds that all phonemic synaesthetes present featural colour patterns, and that features which are acquired the earliest are most likely to influence synaesthesia. Participants' systems were also sensitive to prosodic properties, abstraction and underspecification. This shows that synaesthesia has access to the phonological module and its subunits, providing further evidence that they are psychologically real. This implies that synaesthetes' intuitions can be used to probe the phonological component to provide evidence for linguistic phenomena.

2 Review of the literature

In order for the hypothesis that phonological units influence synaesthesia to be viable, it is necessary for both phonological units and synaesthetic experiences to be psychologically real - an interaction between the two is improbable otherwise. In addition, it is helpful to show that there is precedent for linguistic involvement in synaesthesia and that both synaesthetic correlates and phonological units are acquired within the same time window as this would make it more probable for one to influence the other. This section will start by discussing features, evidence for their psychological reality and theories regarding their acquisition, then discuss synaesthesia, what it is and what is known about its interaction with linguistics.

2.1 Features

2.1.1 Phonemes as bundles of features

In the generative tradition, phonemes are thought of as bundles of features (Distinctive Feature Theory; Jakobson, Trubetzkoy, and Karchevsky 1928). These features have been variably argued to be acoustic or articulatory (Browman & Goldstein, 1992) in nature, a mixture of the two, or simply a way to describe class behaviour and therefore to have no inherent content (Hale & Reiss, 2008).

There is evidence from a wide range of areas that features are real in the mind of the speaker:

- Disorders - aphasias can selectively target consonants or vowels, providing evidence for a [±consonantal] feature (Cotelli et al., 2003; Ferreres et al., 2003);
- Mismatch fields - Phillips, Pellathy & Marantz (2000) contrasted phonemes differing in [±voice] in a modified auditory mismatch paradigm - a mismatch field was elicited in the left hemisphere when the subjects were presented with a [±voice] deviant, showing that the brain has access to subphonemic features.
- Activation - Scharinger et al. (2016), assuming an underspecification approach, performed a same/different task which involved hearing two vowels in sequence and measuring the activation of the superior temporal sulcus (an area found to be important for speech sound processing) - it was found that if the first vowel was specified for feature X and the second was not, there was stronger activation of the STS than in the reverse order, which is a result not only showing the psychological reality of features, but specifically of an underspecification account. In addition, Lawyer and Corina (2014) found evidence in an fMRI-adaptation experiment of areas in the superior and medial temporal lobes which respond selectively to changes in the major feature categories of voicing and place of articulation.
- Localisation of features - Obleser, Lahiri and Eulitz (2004) have produced MEG evidence which suggests that different vowels are localized in the auditory cortex in a way that reflects their feature membership when it comes to place of articulation.

The robust evidence for features suggests that natural classes (sets of phonemes picked out by a feature or set thereof) are also psychologically real, and therefore may be able to play a role in phenomena which appear to work with mental units, such as synaesthesia.

2.1.2 Feature Acquisition

Of relevance to this work is the order of acquisition of features: because synaesthesia is either innate or develops extremely early, one may expect the patterns found in synaesthetic colour correlates to be particularly sensitive to the features that are acquired the earliest, as those which are acquired later may end up outside the sensitive period for the assignment of synaesthetic sensory correlates.

When learning a language, children must learn which of the subphonemic contrasts are relevant for meaning, i.e. which features are relevant to their mother tongue. There is debate on whether this is done by extrapolating meaningful contrasts from fully specified representations (which entails features being binary) or by splitting the phoneme inventory until every phoneme is uniquely identifiable by its feature combination. The first position is untenable, as it has both conceptual and empirical limitations (see Dresher, 2003 for a summary), therefore recent models rely on the second type of mechanism. Such models are either entirely probabilistic in nature (Pierrehumbert, 2003) or rely on an innate hierarchy of features which are progressively employed to subdivide the phonemic space (Dresher, 2009) - this dissertation will assume the latter.

Emergentist models aim to account not only for universality but also for individual variation in the order of acquisition. This is done by positing a hierarchy of features in which a mother node must be acquired before any of the daughter nodes, but sister nodes can be acquired in any order. The feature hierarchy which is usually accepted is that of Jakobson and Halle (1956), shown below.

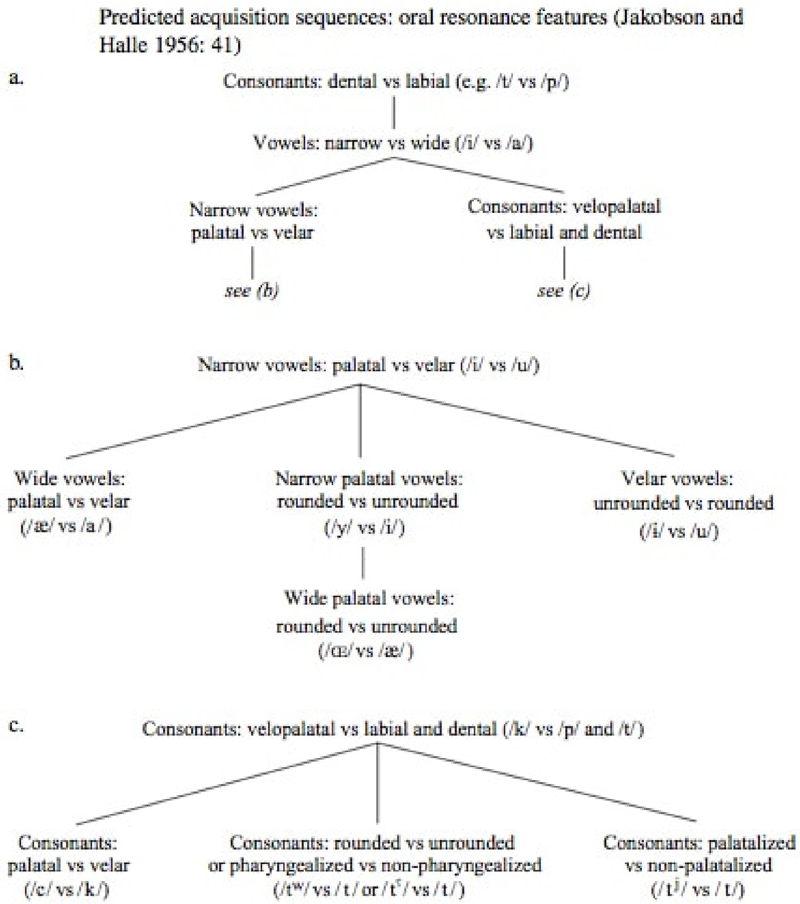

Figure 1: Contrastive Feature Hierarchy - Jakobson & Halle (1956)

Given the above, one would expect any colour correlate patterns to be more likely to be sensitive to the features closest to the root of the feature hierarchy, such as nasality/sonorant/continuant etc., rather than to [ATR], [spread glottis], etc..

2.2 Synaesthesia

2.2.1 Ontology and diagnosis of synaesthesia

Synaesthesia is a neurological condition in which stimulation of one sensory pathway provokes involuntary and automatic experiences in a second sensory or cognitive pathway (Simner, 2007). It occurs in about 1% of the general population (Simner et al, 2006), and may happen between any two senses or cognitive concepts. Cytowic (2002) describes the necessary conditions for experiences to be classified as synaesthetic as the following:

1. Synesthesia is involuntary but elicited.
2. Synesthetic perceptions are spatially extended.
3. Synesthetic percepts are consistent and discrete.
4. Synesthesia is memorable.

5. Synesthesia is emotional.

Synaesthesia is usually tested through Baron-Cohen et al. (test of genuineness which emphasises Cytowic's third criterion as, in its classical version, it consists of eliciting a synaesthetic judgement of the same stimuli a year after the first test which, if congruent, indicates synaesthesia. A different method, also emphasising the third criterion, is a modified Stroop test in which reaction times are measured in conditions where the priming element is either congruent or incongruent with the stimulus: if a synaesthete sees 4 as blue, being presented with orange priming 4 has been shown to result in longer reaction times, while this is not the case in controls (Mills et al., 1999).

2.2.2 Aetiology of synaesthesia

The predisposition for synaesthesia can be innate (Baron-Cohen et al., 1996) or acquired through extensive training or brain lesions (Colizoli et al, 2012). The age of onset of synaesthesia is universally agreed to be very early in the development of the child, although there are disagreements on exactly how early: according to the neonatal hypothesis (Maurer, 1993), everyone is born with synaesthesia and the condition slowly disappears in the majority of the population as the individual ages. Wagner et al. (2011) have shown that the presence in the environment of particular shapes influenced colour preferences in 2- and 3-month olds, but this effect disappears in most 8-month olds and adults, which does appear indicative of early synaesthetic experiences.

2.2.3 Phonemic synaesthesia

Little is known about phonemic synaesthesia, other than the fact that it does exist. Simner (2006) suggests that the literature has been using the term to indicate synaesthesia triggered by linguistic auditory rather than graphemic stimuli. If the trigger for the synaesthetic experience is auditory and the effect is not recreated when the individual is reading, the pathway of word processing which activates synaesthesia must be phonemic rather than graphemic, suggesting that this type of synaesthesia involves phonemes. This possibility will be tested in this thesis.

2.3 Bridging the two literatures - linguistic factors in synaesthesia

Linguistic and paralinguistic factors have been shown to influence a synaesthete's colouring of words. There are lexical synaesthetes, who have a colour for the whole word but not for single letters or sounds (e.g. EP in Baron-Cohen et al., 1987), who presumably have a binding to the meaning rather than the form of the word. The rest of the synaesthetic population is sensitive to the internal structure of the word. Previous studies (Baron-Cohen et al., 1993; Cytowic, 1989; Marks, 1975; Ward et al., 2005) have emphasized the importance of the serial order of segments in deciding word colouring: words seem to get their colour from a) their initial letter; b) their initial vowel or c) a combination of the two.

Simner et al. (2006) suggest that it is not always the initial vowel/consonant which colours the word. They found that, at least for some synaesthetes, prosody matters: colouring will depend on the stressed syllable first consonant/vowel. Previous research used nouns in synaesthesia experiments, 90% of which have initial stress in English, thus masking the importance of prosody. Therefore, linguistic information such as prosody is available to synaesthesia.

Moreover, Mankin et al. (2016) show that morphological information is relevant to synaesthesia by looking at the processing of compound words, and frequency has also been shown to have an effect (Beeli et al. 2007). It is therefore reasonable to hypothesize that phonology and features may be able to play roles in synaesthetic systems.

3 Methodology

3.1 Preface to the experiment

The discussion above has shown that synaesthesia, phonemes and features are psychologically real and have even been localised in the brain and, furthermore, previous experiments show that linguistic information informs synaesthetic correlates. It therefore seems reasonable to propose that features and natural classes too may have a role to play in the colour correlates of phonemes. This is the hypothesis I shall investigate. In this section I will address the choices made in designing the experiment, and then go on to consider the implementation of the experiment itself.

3.1.1 In defense of single-subject experiments

There are widespread assumptions in the scientific community that experiments focussing on single subjects are low in generality or questionable due to the lack of statistical tests (Dywan & Segalowitz, 1986; McCoy & Pany, 1986; Barlow, Hayes & Nelson, 1984; Gelb, 1997 - amongst many others), and a tendency to aggregate subjects or even experiments in order to attempt to reach more and more general conclusions. However, such attempts inevitably mask variability, potentially leading to inaccurate conclusions.

Single-subject experiments have been accused of low generality. Before addressing this, it is crucial to recognise that even multi-subject experiments are often low in generality due to the difficulties in achieving true random sampling. An objection to this criticism from believers in Universal Grammar (Chomsky, 1957, 1965, 1970, 1981, 1993) or a similar form of universal language mechanism is based on the (theory-specific) assumption that said faculty is a property of all humans, and so even a biased sample can yield valid generalisations. Therefore, if generalisations can be made from a biased sample, single-subject results should be equally generalisable. (This is already applied in the field of syntax, which are often based on the grammaticality judgements of very few individuals).

Generalisations are obtained by carefully analysing all factors involved in a process and isolating causal relationships, and at a later stage ensuring that the findings can be replicated across operationally different but theoretically identical experiments, with the aim of ascertaining the exact domain of the generalisation. In the words of Thorngate (1986), “one must find out what each person does in particular, then determine what (if anything) these particulars have in common”. It is not the experiments themselves which can be generalised, but aspects of them - therefore the issue is whether the relationships found in single-subject experiments are in any way less generalisation-worthy than those found in multi-subject experiments, which is an empirical issue that requires further attention and therefore cannot be treated as an assumption.

Multi-subject experiments use statistical tests in order to control for variability, therefore research that foregoes the use of statistical tests must simply find other ways of controlling for variability, such as experimental control (Michael, 1974). Experimental control is found in multi-subject research in the form of randomly assigning subjects to different conditions, ensuring the conditions are kept constant and then averaging the data across subjects. In a single-subject experiment, this is equivalent to directly eliminating sources of variability until the effects are apparent for each subject. If done properly, this makes statistical significance tests unnecessary.

Moreover, Gallistel et al. (2004) show the risk of averaging across participants: looking at learning curves describing classical conditioning, averaging across participants makes it appear that this happens according to a gradual slope - however, a closer look at individual results shows that each individual shows a sudden increase in response to conditioning, asymptotic levels of performance differ greatly between subject, latency length varies greatly across subjects, there is no correlation between the length of the latency period and the performance asymptote, and the performance may decrease after plateauing. All of this is apparent when looking at individual results separately, and the human brain can easily make the generalisations described above, but the results are obscured when the analyst averages over sets of individuals. The above illustrates how statistics doesn't necessarily make results clearer or more valid - and conversely, how a lack of statistic analysis doesn't necessarily mean the raw data cannot show interesting patterns that wouldn't show up otherwise.

And finally, of specific relevance to linguistics and this study, each individual effectively has their own linguistic and synaesthetic system, which may differ substantially from those of their peers. Therefore given that the aim is to investigate individual phonological systems, it appears sensible to look at each system on its own merits rather than attempting to average across participants. Finding several, albeit differing, colour patterns in the sample should be taken as evidence in favour of the hypothesis even when statistical tests are lacking.

3.2 Participants

Three participants were deemed fit for the study, a number which led to the decision of performing in-depth analyses of individual systems instead of doing a quantitative analysis over all participants, especially in light of the significant variability between synaesthetic systems discussed above. An unfiltered sample of prospective participants was first gathered through a variety of strategies including posting on Facebook™, word of mouth, and putting up leaflets in faculties that were statistically more likely to contain synaesthetes (Maths, Music, MML, Psychology, as per Ward et al. 2008). The challenge of explaining what phonemes and synaesthesia are was addressed by exemplifying the experience of a phonemic/graphemic synaesthete ("does P have the colour green? Does the word "refrigerator" have a purple hue?"), and then briefly explaining what synaesthesia is. A phone interview was arranged with all respondents in order to flesh out their specific system with the aim of isolating phonemic synaesthetes.

3.3 Phone interview

Each phone interview lasted about 10 minutes. Participants ensured they were in a quiet area with good reception so as to allow phonetic details to carry through as well as possible. Sibilants were avoided in test words as the higher frequencies that characterise them fall outside of the telephone's bandwidth, making them harder to identify (Coe, 1995). The choice of a phone interview rather than a written questionnaire was designed to minimize the possibility of graphemic interference, and to decrease the risk of not individuating a phonemic synaesthete in case the individual happened to have both graphemic and phonemic synaesthesia.

Participants were asked their age and L1, to identify whether English was their first language and/or how long they had been learning it. The purpose of this was to make sure their feature specifications for English phonemes converged with those of native speakers (predicted to be possible under emergentist assumptions about feature acquisition, section 2.1.3). L2 English speakers who could not pass for native speakers were excluded from participating in the experiment. RG has English as their L1, MG is bilingual in French and English and MC passes as a native speaker even though their L1s are Italian and Romanian.

They were invited to describe their synaesthesia - this was used to ascertain whether the prospective participant knew what synaesthesia was and was not mistaking it for phenomena of similar symptoms such as olfactory memory (when presented with e.g. BENCH, the subject associates benches with an event in which chicken was eaten while sitting on a bench, and therefore associating BENCH with the taste of chicken) or ideasthesia (e.g the Kiki/Bouba effect (Ramachandran & Hubbard, 2001), the simple association of two concepts due to external factors such as, in the above case, the "sharper" or "blunter" sound of the word, a case which is a type of phonaesthesia). They were then questioned about details of their experience to ensure it aligned with Cytowicz's diagnostic criteria for synaesthesia (2002). In addition, if the age of onset of synaesthesia was in adulthood the participant was discarded, as late onset is linked to potentially interfering factors such as temporal lobe lesions, schizophrenia and medication for other conditions, and was therefore not useful for the investigation of phonemic features which would have been acquired in early childhood. If the prospective participant was ascertained to a) have a good enough level of English and b) meet the diagnostic criteria, the phone call would proceed.

The following questions were designed to weed out various types of synaesthesia that would not be useful for the experiment, and to identify any complications in the individual's system which would potentially interfere with the results, in order to adjust for these in the testing. To eliminate the possibility of their synaesthesia being purely semantic or lexical, subjects were presented with

- Words from multiple grammatical classes - semantic synaesthetes are expected to lack a colour for function words as these have no fixed semantic properties;
- Abstract words - a system which colours abstract words is a sign that its colour associations pertain to the word's form, not to the real-world object its meaning picks out;
- Emotionally charged words - as some synaesthetes associate words which elicit emotion with a colour as a result of associating the emotion itself with a colour - usually words that cluster together emotionally, such as family words (MOTHER-FATHER-CHILD), would have the same colour for this kind of individual. Individuals with this characteristic were not discarded, but this was noted in their profile. Of our three participants, only MG had

this.

- Homonym pairs - these were presented in pairs with a priming word, e.g. BURGER-CHIP; WOOD-CHIP : CAR-PARK; BENCH-PARK. If the colour given for the word remained invariant regardless of the priming word it meant that the synaesthesia was dependent on the string and not on the semantic properties of the word.
- (Near-)homophone pairs - to minimise the likelihood of graphemic interference, subjects were asked for colour correlates for FARM - PHARMACY, in which the string at the beginning is phonologically identical but graphemically distinct. Having the same colour for both was a sign of phonemic synaesthesia.
- Real-world coloured items - GRASS, SKY, PIG to ensure that the subject was not simply associating the real-world colour of an item to its label.

After this, the prospective participant was asked whether single phonemes elicited a synaesthetic response in order to direct them to the right link. Only MG said that consonant sounds on their own did not elicit a correlate, and was tested to make sure both vowels and consonants elicited synaesthetic responses (either by producing a sound and asking for the correlate or presenting the participant with words beginning with the sound of interest). Participants who passed all crucial stages were then sent the experiment, and a note was made on their profile about any particular characteristics of their systems.

3.4 Experiment

All participants completed subpart A, which contained vowel sounds in isolation. Those who reported isolated phonemes eliciting colour (MC, RG) completed subpart B, which consisted of consonants in the template [ɑː], the rest (MG) complete subpart C, for which the words used are in Table 1. Where possible, the prompt eliciting the same variable for subparts B and C was kept identical, e.g. to elicit /k/, [kɑː] was used in B and CAR - also [kɑː] - in C. The stimuli were presented in an auditory modality to minimise grapheme interference.

The tests had the interface in Figure 2: by pressing PLAY the participant would hear the relevant sound and then be asked to select the colour on the grid (18x19 squares, hue as the Y axis and luminosity/vibrance as the X axis) that best corresponded to the stimulus. If the sound played triggered no colour correlate, the participant was instructed to click the bar between the grid and the PLAY button, where the selected colour normally appeared.

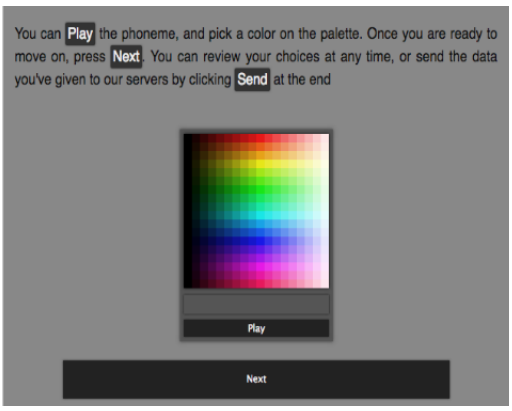

Figure 2: The experiment's subject interface

The stimulus could be replayed and the correlate could be changed at any point before submission of results. Participants were asked to do the experiment with headphones in a quiet environment. The sound prompts were recorded in the phonetics laboratory of the university, and all sound files thus produced were stored in uncompressed .wav format.

Figure 3: The experiment's admin interface

The colours were then recorded in a separate admin page in the format in Figure 3 - the numbers on the colours correspond to [hue:luminosity] - where the experimenter could see the phoneme it corresponded to by hovering the cursor over it.

In addition, RG and MC both underwent an interview session designed to understand the subtleties of their system and to elicit problematic phonemes (section 3.5). Not all phonemes were included in the test, for reasons discussed in the next section.

3.5 Methodological issues and challenges

The greatest challenge for the study was recruiting participants: this was due not only to the limited size of the relevant clinical population, but also to difficulties explaining to laypeople what synaesthesia and phonemes are, in order for them to be able to respond to a request for participants. The solution adopted was to forego mentioning phonemes and synaesthesia altogether, and instead exemplify the phenomenon - "is the word "refrigerator" red? Is the letter p yellow?" - which matched the experience of phoneme-/meaning- and grapheme-colour synaesthetes. Phoneme-colour synaesthetes were then isolated through the phone interview.

Given time and funding circumstances, the above was the best possible method for the aim of the study, although it presents a series of potential problems: first of all it was impossible to ascertain beyond doubt that the person had synaesthesia and not a simple association between words/letters and colours based on e.g. a coloured alphabet they had learned in childhood. A test of genuineness (TOG-R, Asher et al., 2006) would have solved this issue, but the individual must repeatedly take the test overseveral months or even years to ensure consistency, which was not feasible given the time frame for this study. Questions were therefore included in the phone interview (section 3.3) about how long the subject had had synaesthesia and it was noted if their responses when asked for a colour correlate were not almost immediate (if there is a significant delay between the question and the answer about a synaesthetic correlate, it is safe to assume that the correlation is not automatic, which is required for synaesthesia, and the lack of automatism in turn indicates that the colour correlates are not constant over time), which however entails that the reaction times and consistency over time are not measured in a systematic manner but are simply estimated.

A potential risk was including a predominantly graphemic synaesthete: the phone interview method was chosen to minimise graphemic interference, but it has long been known that orthography plays a role in speech perception (Seidenberg and Tanenhaus, 1979; Donnenwerth-Nolan et al. 1981) so there is potential for a graphemic synaesthete to pass for a phonemic one. To minimise the likelihood of this, potential subjects were asked the colour correlates of PHARMACY and FARM, but it cannot be excluded that <f> and <p(h)> may incidentally have the same colour correlate (although the statistical likelihood of this is low), leading to a grapheme synaesthete being incorrectly identified as a phonemic one. A research project with more time and funds may wish to design a test with visual distractors (either in the form of black and white pictures or unrelated words) to minimise graphemic interference.

As for the design of the experiment, some phonemes could not be included due to various potential difficulties in eliciting their colour correlates, which may result in skewed data. However, in many cases, the excluded phonemes would have provided crucial pieces of data to understand what patterns were really at work, so any further research in this area will have to work around this difficulty. Below are the phonemes that have been left out and why the choice was made:

- Glides - phonetically, glides and high vowels have conspicuous overlap, as their main difference is a phonological one: their [±syllabic] feature. Therefore given that the experimental setup involved eliciting segments in a [_ɑː] template for consonants and in a [_] template for vowels, glides could have potentially been interpreted as the onset of a diphthong and not as segments in their own right, or as vowels. This is unfortunate, given that finding a difference in colour between /j/ and /i/; and /w/ and /u/ would have been synaesthetic evidence for the existence of [±syllabic], as except for this feature the two members of each pair are identical phonetically and featurally[1] A possible solution would perhaps involve having a graphemic representation along with the auditory elicitation exclusively for high vowels and glides (<ja> would prime the subject to think of a consonant while <ia> would invite them to think of the same segment as a diphthong) but this would also increase the likelihood of graphemic interference.

- Central vowels[2]: these have been excluded because, when presented in isolation, they are too easily mistaken for a reduced version of one of the peripheral vowels. Presenting them in words (e.g. for /ɪ/, when the subject was presented with the word /ɪnstɪnkt/ there would have been increased likelihood of the colour of the first consonant /n/ giving colour to the word, as /ɪ/ might be too "weak" to impose its colour over a comparatively stronger consonant (one participant reported this phenomenon)

- /ŋ/ - as this phoneme only appears in codas in English and synaesthesia works mostly in a linear fashion when ascribing colour to a string of sounds (as outlined in section 2.2), the risk with using the string /aŋ/ to elicit its colour would have been likely to return the correlate for /a/. The other option would have been presenting the sound on its own, as it is after all a sonorant which can be pronounced in isolation - the issue with this is that, because of the possibility of graphemic interference combined with the unusualness of being presented with the sound in a non-coda position, the untrained ear could have analysed it as /n/ and returned the correlate for it instead of /ŋ/.

  Moreover, a statistical analysis of the results would have corroborated the findings - a modified version of a phylogenetic regression (Grafen, 1989) could have been applied to the data in order to dismiss phonetic similarity as a source of interference in the patterns below that appear featural in nature. The reason why it was not applied here is its elevated complexity

[1] This is not universally agreed upon, but see Padgett, 2008 for evidence in favour of this position.

[2] These were however elicited from RG and MC, ensuring it was clear to them what segment was being elicited and testing repeatedly for the same segment with words that represent the sound with different graphemes, to minimise the risk of graphemic interference. This was not possible with MG due to geographical constraints

4 Results

Given the discussion in section 2 about the ages of onset and stabilisation of phonological and synaesthetic systems, I provisionally formulate the Photograph Effect Hypothesis which, when coupled with the feature hierarchy, makes predictions about what features one may expect the colour correlates of a synaesthetic system to be sensitive to.

> **The Photograph Effect Hypothesis (PEH)**
> Because synaesthetic systems are consolidated early in the development of the relevant sensory and cognitive pathways, a coloured synaesthetic systems may provide a fossilised image of the cognitive system it is linked to at the time of consolidation T of the correlates of the synaesthetic system.

The following section isolates interesting patterns found in the system of synaesthetes and attempt to explain them in light of the PEH and the feature hierarchy. When two hues are described as close, this is intended to mean that they are not more than 90° apart in the colour wheel in Table 2 (a 4-primes colour wheel, also known as a natural colour wheel, because it better represents relationships between colours as we perceive them), while when they are described as very close, they are not more than 45° apart.

4.1 Consonants

4.1.1 MC [3]

	MANNER		VOICING	PLACE						
				Bilabial	Labiodental	Interdental	Alveolar	Palatal	Velar	Glottal
Obstruent	Stop		Voiceless	p			t		k	ʔ
			Voiced	b			d		g	
	Fricative		Voiceless		f	θ	s	ʃ		h
			Voiced		v	ð	z	ʒ		
	Affricate		Voiceless					t͡ʃ		
			Voiced					d͡ʒ		
Sonorant	Nasal		Voiced	m			n		ŋ	
	Liquid	Lateral	Voiced				l			
		Rhotic	Voiced					r		
	Glide		Voiced	w				j	w	

Figure 4: MC's consonantal colour correlates

Several patterns appear: fricatives are either purple or blue, which are very close on the spectrum. Moreover there is a general division between voiced ones, which are of purple

[3] ŋ and ʔ were included for MC as a more detailed study was performed on this participant, and as such the obstacles to obtaining these data points were able to be overcome. See section 4.3.

hue, and voiceless ones, which are blue in hue. All nasals are red, differing only in luminosity. Both liquids are shades of green or yellow (close on the spectrum) and have low luminosity.

All velar segments are of very low luminosity, and cluster around the red-brown part of the spectrum. All coronal segments are blue or green (colours which are close on the spectrum) except for /n/, suggesting perhaps a) a colour correlate for place features and b) that nasality ranks higher in a feature hierarchy (as predicted by feature hierarchy models such as Jakobson Halle, 1956) and its signature colour correlate will therefore trump that of the place of articulation.

The pattern of affricates here is typical of all other systems: their colour correlate appears to combine those of the corresponding stop and fricative. This gradient phenomenon is not necessarily expected if one posits that all processes which are phonological are discrete, as this would mean that affricates are assigned colour correlates based on their phonetic composition rather than their natural class membership. This can be accounted for when considering the fact that affricates are nowadays analysed as a cohesive block which however maintains separate feature specifications for its components, which explains why their colour correlate is a mixture of that of their components' correlates.

Finally, /h/ is very light, almost white, in all subjects - the reason for this may be that it is specified for very few features(e.g.no place specification (Schluter et al, 2016)) , which may result in lack of vibrant coloration. A similar thing seems to be occurring for the glottal stop, which is both black and white; perhaps, because it can stand for all voiceless stops, it is perceived as inherently underspecified (as shown in section 4.3).

4.1.2 RG

	MANNER		VOICING	PLACE						
				Bilabial	Labiodental	Interdental	Alveolar	Palatal	Velar	Glottal
Obstruent	Stop		Voiceless	p			t		k	
			Voiced	b			d		g	
	Fricative		Voiceless		f	θ	s	ʃ		h
			Voiced		v	ð	z	ʒ		
	Affricate		Voiceless					t͡ʃ		
			Voiced					d͡ʒ		
Sonorant	Nasal		Voiced	m			n			
	Liquid	Lateral	Voiced				l			
		Rhotic	Voiced					r		
	Glide		Voiced	w				j	w	

Figure 5: RG's consonantal colour correlates

RG shows a pattern for stops similar to that found in MC's fricatives: all stops are pink-blue, which are close in the spectrum, with the voiceless ones being pink and the voiced ones being blue. In addition, there seems to be a correlation with place of articulation within stops: the farther back the place, the lower the luminosity.

Apart from voiceless stops, the subject shows a general division on the basis of place of articulation in all segments, with coronal segments being green-blue and dorsal ones being yellow, purple and brown[4], with the exception of nasals (which in MC's system also have their colour correlate overinding other potential colours, suggesting that nasality ranks high in the contrastive feature hierarchy) and of /z/: while the cause for the latter oddity cannot be explained simply in featural terms, it is worth noting that /z/ also behaves unexpectedly for MG, so there may be a principled reason for this departure from expectation. The unexpected behavior of nasality may be because it ranks significantly higher in the feature hierarchy than place of articulation, and therefore the colour correlate for it (red-pink, very close colours in the spectrum) trumps all others.

As with the other phonemic subjects, RG appears to consider affricates as composite segments, and adheres to the pattern of /h/ being significantly different from all other segments. However, s/he uses a different strategy from MC and MG to signify underspecification, colouring the segment a low-luminosity mud brown hue. In colour theory, brown is obtained by mixing complementary colours, or by mixing any secondary colour with a subsequent one (Place Madry, 1989) which from a derivational perspective makes brown an indeterminate colour, as it could come from any number of combinations - it appears therefore quite natural for a segment assumed to be underspecified for many features to be coloured brown.

[4]Brown does not appear in the colour wheel for reasons to do with colour theory. However, there are different ways to obtain it, and one of these is by mixing purple and yellow (Place Mardy, 1989). This fact gives more cohesiveness to the class of dorsal segments.

4.1.3 MG

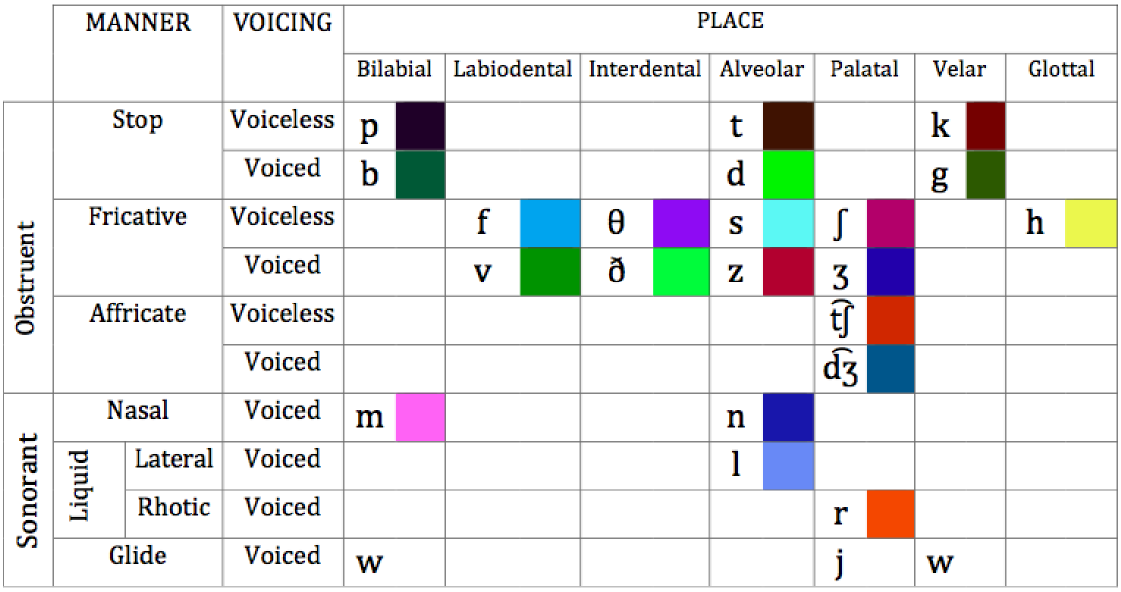

Figure 6: MG's consonantal colour correlates

MG presents a similar pattern to RG for stops, with voiceless ones being black/brown and voiced ones being green. It could be argued that stops also present place feature correlates in the form of luminosity, with coronals being the most luminous and labials being the least. This is reminiscent of Pater Werle's (2001) findings that their child subject had a feature hierarchy for place of articulation in assimilatory processes, coronal > dorsal > labial. This system might have been shared by MG in childhood and during the solidification of synaesthetic judgements this might have been an influence.

Moreover, apart from /h/ which is discussed below, MG presents a pattern on the basis of voicing in the fricatives: voiceless ones are blue-purple, very close on the spectrum, and voiced ones are green-blue, also very close on the spectrum, with the exception of /z/ which is red - this is identical to RG's pattern for place of articulation. It appears that there be something about /z/, in its phonetics or phonology, which makes it an outlier;. this deserves more attention in future work.

MG also shares the patterns of "light /h/" and the affricates seen as composite of the relevant stops and fricatives with all subjects.

4.2 Vowels

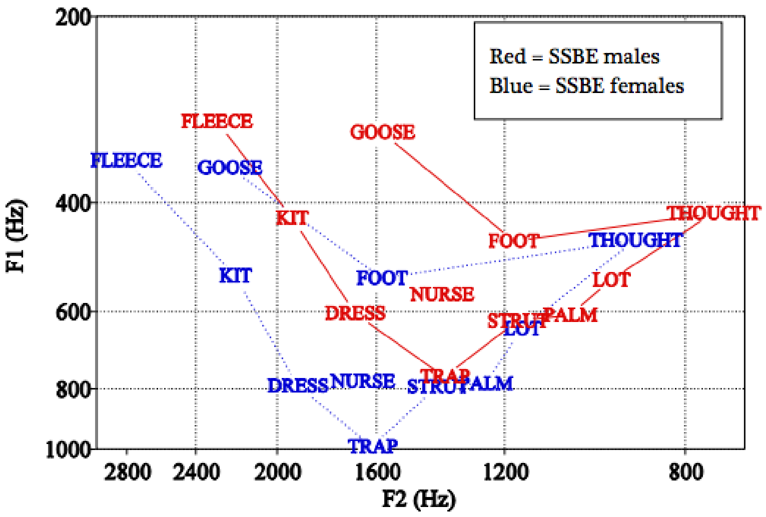

Figure 7: Perceptual similarity of SSBE vowels (Williams, 2013)

The figure above represents perceptual similarity of SSBE vowels by sex (Williams, 2013). Note the compression of the higher frequencies on the scale, meant to better represent perceptual similarity. The figure will be invoked when phonetic similarity could play a role in the proposed patterns.

4.2.1 MC[5]

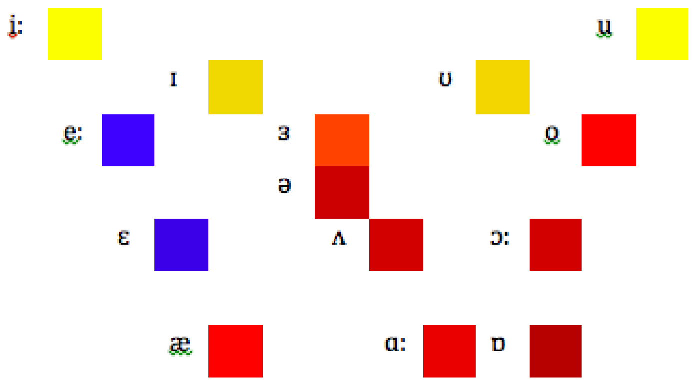

Figure 8: Perceptual similarity of SSBE vowels (Williams, 2013)

This system shows both height and backness patterns. The patterns are indicative of a hierarchy of features, with height ranking above backness: [-low, +high] vowels are yellow and [+low, -high] vowels are red, regardless of backness specifications.

[5]MC's system is treated more thoroughly in section 4.3

The picture painted by the mid and central vowels is interesting: note first how the only difference between the peripheral mid vowels of the same height is a slight one of luminosity. This could be ascribed to the late acquisition of the [ATR] feature (Jakobson Halle, 1956 and all subsequent work on feature hierarchies). This is also borne out by the behaviour of [ATR] pairs in the high vowels. The general pattern seems to be that [+ATR] vowels are marginally more luminous than their [-ATR] counterparts. The data is better understood in the light of the fact that, in SSBE, /u:/ and /o:/ are just about equidistant from /ʊ/ perceptually as is shown in the table at the start of the subsection. And yet, /ʊ/ does not colour itself as a mixture of the two and is quite clearly closer to /u:/ in hue. So assuming the PEH, one can suggest that by the time of consolidation of the system, height contrasts had been acquired but [ATR] contrasts had not, and therefore /ʊ/, as a high vowel, patterns with /u:/. This would explain why the difference between [ATR] pairs is consistently a slight and gradient one: because it's phonetic, not phonological.

Moreover, note how // is red and patterns with the back vowels even though it is phonetically central, which is a clear indicator that such systems are predominantly phonological rather than phonetic.

This system is consistent with the predictions made by the conjunction of the PEH with a feature hierarchy: height contrasts should be the first to be acquired (represented in colour by the fact that colour correlates for height outrank those for other features i.e. /i/ has the correlate for [+high] and not that for [+front]), then a backness contrast, and only after this the contrast in height between mid vowels (featurally an ATR contrast). The coloured system shows a height contrast and a general backness contrast, but ATR contrasts are not at all as clearly marked as other featural contrasts, as they were phonetic at the time of system consolidation.

4.2.2 RG

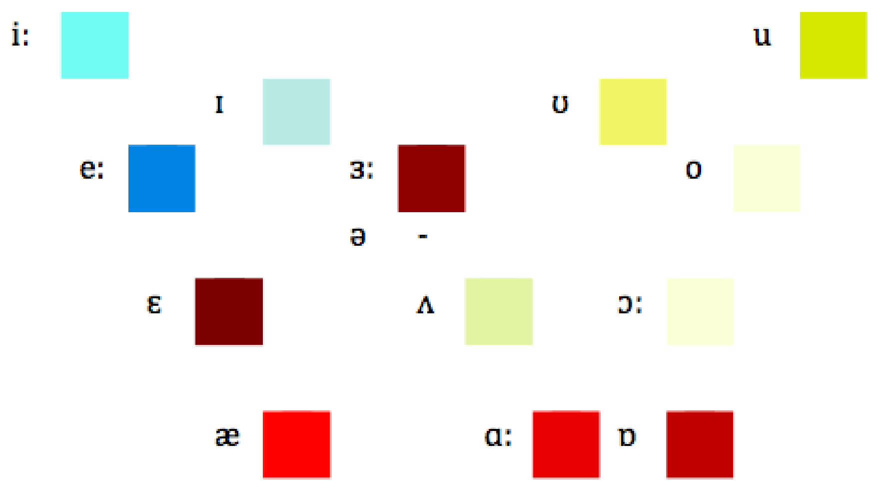

Figure 9: RG's vocalic colour correlates

Low vowels are the same hue, suggesting a correlate for [+low] while high vowels do not share a hue but all have very high luminosity. This is in line with acquisition predictions about a height distinction being acquired before a backness one, as all low vowels are the same hue regardless of backness, as a result of height outranking backness in the contrastive feature hierarchy, and therefore also outranking it when imposing a feature correlate. An argument may be made for blue correlating with frontness and yellow with backness - that is what differentiates /iː, eː, i/ from /u, ʊ, ʌ/[6]. A differentiation in frontness between high (but not low) vowels which is predicted as the second crucial differentiation by Jakobson Halle (1956).

The peripheral mid vowels are not differentiated: in the back, mid vowels appear to have differentiated from peripheral vowels but not between each other, while in the front they don't seem to have differentiated from the peripheral vowels - based on colour and on the assumption of PEH, it seems that /e:/ may have been treated as an allophone of /i:/ and /ɛ/ as an allophone of /ae/ at the time of consolidation of the system, before the category of mid vowels had been acquired and young RG had to make choices about how to divide the phonetic space. This cannot be simply a phonetic effect, as for example /o:/ is perceptually much closer to /u:/ than /e:/ to /i:/(Figure 7), and yet the first pair does not show this effect. RG reports that /ə/ in isolation has many colours, and it is only when it is put in context that it can be assigned the colour of the phoneme it alternates with: if presented with MOUNTAIN [ˈmaʊntən], /ə/ took the colour of /ɛ/, which is the non-reduced phoneme in that context as shown by MOUNTAINOUS [ˈmaʊntɛnəs][7] - this is almost identical to MC's treatment of the segment.

Assuming PEH, the situation of the ATR feature is in line with that of MC - it was yet to be acquired at the time of consolidation. While looking at /iː/-/ɪ/ and /uː/-/ʊ/ one may perhaps wish to suggest that ATR is expressed by differences in shade, this is markedly not true when examining the mid vowels, which form two [ATR] pairs but show either no distinction at all in colour such as the back mid vowels, or are of completely different hues such as the front mid vowels. As in MC, this suggests that [ATR] was not yet acquired at the time of synaesthetic consolidation, and therefore colour correlates to ATR pairs were assigned according to other factors.

The treatment of /ɜː/ is unexpected as it does not seem to be part of a pattern which can be expressed in terms of features. Its acquisition as a phoneme is rarely treated in the literature, but as a mid central vowel one can assume it is mastered quite late. It may be that at the time of consolidation RG was treating it as an allophone of /e/ or /æ/ (as discussed before, the two may not have been differentiated at time T).

4.2.3 MG

The pattern here seems to be primarily one of backness: vowels with the same height have the same hue, but the [+back] members of each pair are darker shades. The exception to this are the high vowels, which are two different colours - this is in line with JH's prediction that the second distinction to be acquired is one of frontness between the high vowels. According to Jakobson Halle (1956)'s hierarchy, the first distinction to be acquired is a basic one of height, which this system partly represents as both low vowels being red, but has no overt correlate for the high vowels. It therefore seems that MG's

[6]and potentially /o, ɔː/ if one recognises the PEH, which are in the process of differentiating themselves as mid vowels, hence the less vibrant shade compared to the high back vowels

[7]in RG's idiolect. Moreover, note the lack of graphemic interference as /ə/ is neither the colour of /ɑː/ nor /iː/

system has started consolidating after the second differentiation (and so has RG's, as shown by the high vowels being two different colours; MC's started consolidating after the first step, as their high vowels are identical in colour).

The situation in the mid vowels is unlike that of previous subjects, as here they are differentiated from both the relevant peripheral vowels (unlike in RG) and within ATR pairs, which is an unexpected configuration. The PEH would lead one to suggest an early acquisition of mid vowels and ATR distinctions. As MG was raised bilingual in English and French, the early acquisition of such features may be connected with the comparatively more complex mid vowel system in French (they have extra roundness and nasality contrasts), which might have led the subject to acquire ATR distinctions early as to impose more order on a complex system.

4.3 A closer look at MC's system

In order to show that the phenomenon is in fact a phonological one, and that it has many interesting properties that can most accurately be modelled by reference to phonology (and, more specifically, a theory with some level of abstraction), MC's system was analysed in more detail in an interview-type set-up. Below are some of the relevant findings, which are not meant to be taken as findings in themselves but as directions for future research.

4.3.1 Differing influences of phonology and phonetics

First of all, MC's synaesthesia is quite clearly not reliant on graphemes: for example the vowel /i:/ of QUAY, NINA, KEY and FEE was constant in colour regardless of its spelling, the vowels written as <a> in CANADA (/kænədɑː/ in MC's idiolect) and ALASKA (/əlæskɑː/ for MC) were coloured with different shades of red according to their phonemic membership, and there was no trace of /r/ colour in words that have unpronounced orthographic /r/ in codas, such as NORTH (/nɔːθ/). It is therefore unlikely that any of the effects below are caused by graphemic interference.

While MC's system is mainly phonological, it reveals some phonetic effects, which tend to manifest themselves as partial or gradient influences on the colour of the phoneme taken in isolation, and which are therefore visually distinct from phonological processes, which are more drastic and discrete (as discussed in 4.2). For example, the context in which a phoneme is found will marginally influence its colour as a result of what MC describes as "spreading" of a colour onto an adjacent one. However, there are never drastic changes in the core colour of a phoneme: if a consonant is followed by /i/, which is yellow in MC's system, it is slightly brighter than when followed by /ɑ/, which is dark red, in which case it would be slightly darker.

Moreover, phonetically nasal vowels such as those in TIN or MEN were reported to start off coloured as if found in isolation and then gain streaks of red (the colour of nasality for MC) as the pronunciation of /n/ approached. This mirrors spectral effects of nasality in these words: the vowel starts with its normal formant configuration and amplitudes of lower formants are dampened progressively towards the coda (Hawkins & Stevens, 1985). One might predict that synaesthetes who speak languages such as Polish, in which vowel nasality is phonemically contrastive, will not have nasality as a gradient effect but as a more discrete one.

Phonological effects, on the other hand, are more drastic and discrete in nature. To test this, I recorded the string [ata:] (with VOT of roughly 70ms) in Audacity and derived 9 tokens from it in which the VOT was reduced by 20, 30, 40, 45, 50, 55, 60, 65 and 70 ms. These were then presented in random order and MC was asked to name the colour elicited. There appeared to be a clear categorical perception effect, translating in an abrupt colour change between the tokens with 0 and 5 ms VOT rather than a gradual transition.

The treatment of /h/ is noteworthy when attempting to determine whether the system is based on phonology or phonetics: phonetically, all allophones of English /h/ are voiceless versions of an adjacent vowel (Castelo, 1964), therefore if MC's synaesthesia were based on phonetic rather than phonemic properties one would expect the initial sounds of HEAD and HAT to be coloured differently. MC's system however treats /h/ as a phoneme on par with the others, by giving it its own consistent colour correlate.

Moreover, while some phonetic effects do influence synaesthetic percepts, others do not: as discussed in more detail below, MC associates "heaviness" with phonemic length but not with phonetic length - the vowel in the pair BAD-BAT is of equal "heaviness", even though the first one is phonetically longer due to pre-voice lengthening. The above suggests that in a phonologically synaesthetic system, phonological processes have an impact while phonetic ones generally do not, and when they do it is minor and/or gradient, as with the nasalization example discussed earlier.

4.3.2 Abstraction and underlying representations

MC also shows in visual form interesting properties of phonology which involve abstraction, and therefore are not predicted to be the case if phonemic synaesthesia is merely a surface phenomenon, such as the perception of vowel reduction and phonological alternations. In isolation, // has a red colouring, patterning with the back vowels although it is phonetically central (further evidence that the phenomenon is phonological rather than phonetic), and when it appears in alternation with a peripheral vowel - e.g. in OPERA, /ɔpərɑ/ in MC's idiolect - the schwa is both schwa-coloured and the colour of the non-reduced vowel it originates from, in the above case red and dark blue. This suggests that in such cases schwa holds both the status as a phoneme in its own right and the colour information of the vowel with which it alternates.

Indeterminacy and the effects priming has on it are also represented visually in MC's system: the behaviour of /ʔ/ illustrates this - when presented with the string [taʔ] without pragmatic or visual context (which entails that neither discourse nor the McGurk effect can be used to overcome indeterminacy), where /ʔ/ could be an allophone of any of the voiceless plosives, MC reports seeing the segment as black and white (and distinct from the colour of any of its possible underlying representations), suggesting that the brain may ascribe such colour to neutralising positions on which it cannot make a judgement. This changed with priming: when the string was primed with the word DANCING, the last segment was yellow, the colour of /p/, and when primed with RUBBISH, it became blue, the colour of /t/. This suggests that this component of the system is phonological in nature and that the system itself is capable of performing abstraction.

MC's synaesthesia is also sensitive to syllabic effects: a consonant between non-identical vowels will only be affected by "spreading" (described above) from a tautosyllabic vowel, which shows that these processes are not just a simple issue of adjacency but instead are constrained by abstract prosodic affiliation. The vowel in MEAN /mi:n/, for example, is perceived as having /n/ colour spreading onto it, while this is not the case with the same vowel in MEANIE /miːnɪ/, in which the /n/ is separated from the root vowel by a syllable boundary. This effect suggests moreover that syllables are psychologically real units as they are able play a part in synaesthesia.

Diphthongs are of interest when looking at synaesthesia as a potential diagnostic tool. MC sees diphthongs as being made up of the colours of their components - this means that it is possible to identify the underlying components of the diphthong based on their coloration. The diphthongs in MOUTH and PRICE are pronounced respectively as [aʊ] and [ɑɪ] in SSBE. Surface-oriented accounts posit that their underlying representation is identical to the surface form, while more abstractive accounts suggest that the underlying representations are in fact /au/ and /ai/ and that the surface form is the result of undershoot. MC reports that regardless of microvariations in pronunciation, the colour correlates for the diphthongs were /a/+/i/ and /a/+/u/, rather than /a/+/ɪ/ or /ʊ/. MC's percepts suggest that the second hypothesis is more likely to be correct.

5 Discussion

To conclude, this pilot study appears to support the hypothesis that distinctive features and phonology in general are able to influence the selection of colour correlates. This can be gleaned from the patterns identified in section 4.1 and 4.2, and most clearly in the detailed account of MC's system in section 4.3. Different individuals' systems do seem to be sensitive to different features but the fact is that they are all, in some way, sensitive to features. Several phenomena exemplify how phonemic synaesthesia is principally influenced by phonology rather than phonetics, which suggests that the term "phonological" should be used in the linguistic sense rather than its current use to signify synaesthesia triggered by auditory stimuli. MC's account in general shows that their system predominantly deals with abstract underlying representations and phonological processes, and that phonetic effects may appear but these are more gradient and less pervasive. The difference between phonetics and phonology in influencing synaesthesia can be seen in the vowel systems, in which the differentiation between ATR pairs is often a gradient one, and therefore possibly phonetic, while other features such as [high] and [back] correlate with drastic and discrete differences in colour percept and therefore may be plausibly attributed to phonological influence.

If this hypothesis is accepted, it can be placed in the context of a theory of the order of acquisition of distinctive features[8]. Different theories will make different predictions on the order of acquisition. As discussed at length in section 4, colour correlates of systems appeared to provide a freeze-frame photograph of the state of the phonological system of

[8]in this particular case, a Toronto school emergentist view (Dresher, 2009), as outlined in section 2.1.3, which claims the existence of language-specific feature hierarchies which partly determine the order of acquisition of features but also leaves scope for individual variation

the individual in its early stages. Features which are predicted to be acquired the earliest are the ones synaesthetic systems appear most sensitive to, such as voicing, fricativity, height and basic divisions in place of articulation. The phenomenon has been named the Photograph Effect Hypothesis (outlined at the beginning of section 4). Thanks to the interplay of feature acquisition order and the early consolidation of synaesthetic systems, one can visually test if predictions made by various theories of feature acquisition orders are correct. RG's system is a good example of this: Jakobson Halle (1956)'s contrastive feature hierarchy predict that the first distinction in a vocalic system is one between low and high vowels, and the second is a differentiation in frontness between high vowels but not low ones. This is borne out in RG's system, in which the low vowels are red, the high front vowels are blue and the high back vowels are yellow.

Because of the lack of research on phonemic synaesthesia it is difficult to place this study in the context of a preexisting field. It can, however, add to what is known about synaesthesia, Distinctive Feature Theory and phonological acquisition: the results confirm that synaesthesia can be sensitive to linguistic information - which was known for morphology (Mankin et al, 2016), but not for phonology. Moreover, the results suggest that synaesthesia is a potential new evidence source for phonology: as it is sensitive to phonological units such as features and prosodic domains,, one can conceivably use intuitions of synaesthetes as visual evidence. And finally, given that the acquisition of synaesthesia ends before that of phonology, the PEH provides a useful tool for testing the order of acquisition of features.

If phonology does influence synaesthesia, this entails that, once the mechanisms have been better understood and a larger sample of individuals with phonemic synaesthesia has been examined, one can use synaesthetic judgements as external evidence for phonological phenomena, as these provide a way of probing an individual's phonological system. For example, it may be possible to use colour correlates to identify natural class membership, the domain of a process (cf MC's evidence for the syllable), phoneme membership in disputed cases, or psycholinguistic phenomena regarding perception, amongst other issues. In a future in which phonological synaesthesia is better understood, one may even be able to do fieldwork thanks to phonological synaesthetes - this is a good prospect especially for languages without an orthographic system, which entails that all first-language speakers who present some form of linguistic synaesthesia will be phonological with no graphemic interference - this will also entail a non-negligible number of phonological synaesthetes, as the subset of these that would grow up to become graphemic synaesthetes (Simner, 2006) won't have any graphemes to acquire, remaining therefore phonological, raising the incidence to phonological synaesthesia in such societies to about 1% of the entire population.

A better understanding of the phenomenon might allow for cross fertilisation between the areas of language and synaesthesia acquisition: at the moment, creating a timeline of development in both areas, especially in the early stages, is a challenge due to practical difficulties in working with children. If the hypothesis that a synaesthetic system is a picture of the phonological system at a given point in time, dialogue between the two fields and joint experiments on the matter may help localise both the onset of synaesthesia and a more precise window of acquisition of certain features.

6 Bibliography

Barlow D. H., Hayes S. C., Nelson R. O. (1984). The Scientist Practitioner: Research and Accountability in Clinical and Educational Settings.

Baron-Cohen, S., Harrison, J., Goldstein, L. H., Wyke, M. (1993). Coloured speech perception: is synaesthesia what happens when modularity breaks down? Perception, 22(4), 419–426.

Baron-Cohen, S., Wyke, M. A., Binnie, C. (1987). Hearing words and seeing colours: An experimental investigation of a case of synaesthesia. Perception. https://doi.org/10.1068/p160761

Baron-Cohen, S. (1996). Is there a normal phase of synaesthesia in development. Psyche, 2(27), 223–228.

Beeli, G., Esslen, M., Jäncke, L. (2007). Frequency Correlates in Grapheme-Color Synaesthesia. Psychological Science, 18(9).

Browman, C. P., Goldstein, L. H. (1992). Articulatory Phonology : An Overview *. Haskins Laboratories Status Report on Speech Research.

Castelo, L. M. (1964). The Phonetic and Phonemic Treatment of English /h/. Phonetica, 11(2), 166–123.

Chomsky, N. (1993). A Minimalist Program for Linguistic Theory.

Chomsky, N. (1970). Remarks on nominalization. In Reading in English Transformational Grammar.

Chomsky, N. (1981). Lectures on Government Binding - The Pisa Lectures.

Chomsky, N. (1957). Syntactic Structures.

Chomsky, N. (n.d.). Aspects of the theory of syntax.

Citowicz R. (2002). Synaesthesia - A union of the senses.

Coe, L. (1995). The telephone and its several inventors : a history. McFarland Co.

Colizoli, O., Murre, J. M. J., Rouw, R. (2012). Pseudo-Synesthesia through reading books with colored letters. PLoS ONE, 7(6).

Donnenwerth-Nolan, S., Tannenhaus, M. K., Seidenberg, M. S. (1981). Multiple code activation in word recognition: evidence from rhyme monitoring. J Exp Psychol Hum Learn., 7(3).

Dresher, B. E. (2003). On the acquistion of contrast. Gala 2003, 1(September), 27–46.

Dresher, B. E. (2009). The Contrastive Hierarchy in Phonology.

Ferreres, A. R., Lopez, C. V, China, N. N. (2003). Phonological alexia with vowel-consonant dissociation in non-word reading. Brain and Language, 84(3), 399–413. https://doi.org/10.101 934X(02)00559-X

Gallistel, C. R., Fairhurst, S., Balsam, P. (2004). The learning curve: implications of a quantitative analysis. Proceedings of the National Academy of Sciences of the United States of America, 101(36), 13124–31.

Gelb, S. A. (1997). The problem of typological thinking in Mental Retardation. Mental Retardation, 35(6).

Grafen, A. (1989). The Phylogenetic Regression. Philosophical Transaction of the Royal Society of London. Series B, Biological, 326(1233), 119–157.

Hale, K., Reiss, C. (2008). The Phonological Enterprise.

Jakobson, R., Halle, M. (1956). Fundamentals of Language. Sciences-New York.

Jakobson, R., Trubetzkoy, Karchevsky. (1928). Quelles sont les méthodes les mieux appropriés à un exposé complet et pratique d'une langue quelconque? Reprinted in Jakobson 1971, 3-6.

Lawyer, L., Corina, D. (2004). An Investigation of Place and Voice Features Using fMRI- Adaptation. J Neurolinguistics, 27(1), 2432–2437.

Mankin, J. L., Thompson, C., Branigan, H. P., Simner, J. (2016). Processing compound words: Evidence from synaesthesia. Cognition.

Marks, L. E. (1975). Marks, L.E. (1975). On colored-hearing synesthesia: Cross-modal translations of sensory dimensions. Psychological Bulletin, 83, 303–331.

Maurer, D. (1993). Neonatal Synesthesia: Implications for the Processing of Speech And Faces. In B. de Boysson-Bardies, S. de Schonen, P. Jusczyk, P. McNeilage, J. Morton (Eds.), Developmental Neurocognition: Speech and Face Processing in the First Year of Life (pp. 109–124). Dordrecht: Springer Netherlands.

Mccoy, K. M., Pany, D. (1986). Summary and Analysis Oral Reading Corrective Feedback Research. The Reading Teacher, 39(6), 548–554.

Michael, J. (1974). Research : Mixed Blessing or Curse ?', 4(4), 647–653.

Mills, C. B. (1999). Digit Synaesthesia: A Case Study Using A Stroop-Type Test. Cognitive Neuropsychology, 16(2), 181–191.

Obleser, J., Lahiri, A., Eulitz, C. (2004). Magnetic Brain Response Mirrors Extraction of Phonological Features from Spoken Vowels.

Padgett, J. (2008). Glides , Vowels , and Features. Lingua, 118(1), 1937–1955.

Pater, J., Werle, A. (2001). Typology and Variation in Child Consonant Harmony.

Phillips, C., Pellathy, T., Marantz, A. (2000). Phonological Feature Representations in Auditory Cortex.

Pierrehumbert, J. B. (2003). Phonetic diversity, statistical learning, and acquisition of phonology. Language and Speech, 46(2–3), 115–154.

Place, S., Madry, B. R. (1989). The Art and Science of Professional Makeup.

Ramachandran, S., Hubbard, E. M. (2001). Synaesthesia — A Window Into Perception, Thought and Language, (12), 3–34.

Scharinger, M., Domahs, U., Klein, E., Domahs, F. (2016). Mental representations of vowel features asymmetrically modulate activity in superior temporal sulcus. Brain and Language, 163, 42–49.

Schluter, K., Politzer-Ahles, S., Almeida, D. (2016). No place for /h/: an ERP investigation of English fricative place features. Lang Cogn Neurosci, 31(6), 728–740.

Seidenberg, M. S., Tanenhaus, M. K. (1979). Orthographic effects on rhyme monitoring. Journal of Experimental Psychology: Human Learning and Memory, 5(6), 546–554.

Simner, J. (2006). Beyond perception: synaesthesia as a psycholinguistic phenomenon. Trends in Cognitive Sciences.

Simner, J., Glover, L., Mowat, A. (2006). Linguistic determinants of word colouring in grapheme-colour synaesthesia. Cortex.

Thorngate, W. (1986). The production, detection, and explanation of behavioral patterns. In The Individual Subject and Scientific Psychology.

Wagner, K., Dobkins, K. R. (2011). Synaesthetic associations decrease during infancy. Psychological Science, 22(8), 1067–72.

Ward, J., Simner, J. (2005). Is synaesthesia an X-linked dominant trait with lethality in males? Perception.

Ward, J., Thompson-Lake, D., Ely, R., Kaminski, F. (2008). Synaesthesia, creativity and art: what is the link? British Journal of Psychology (London, England : 1953), 99(Pt 1), 127–141.

å

7 Appendix

p: par	v: vow	dʒ: jar
b: bar	θ: thin	tʃ: char
t: tar	ð: this	h: he
d: do	s: see	m: me
k: car	z: zoo	n: knee
g: go	ʃ: shoe	r: read
f: fee	ʒ: gee	l: loo

Figure 10: Table 1

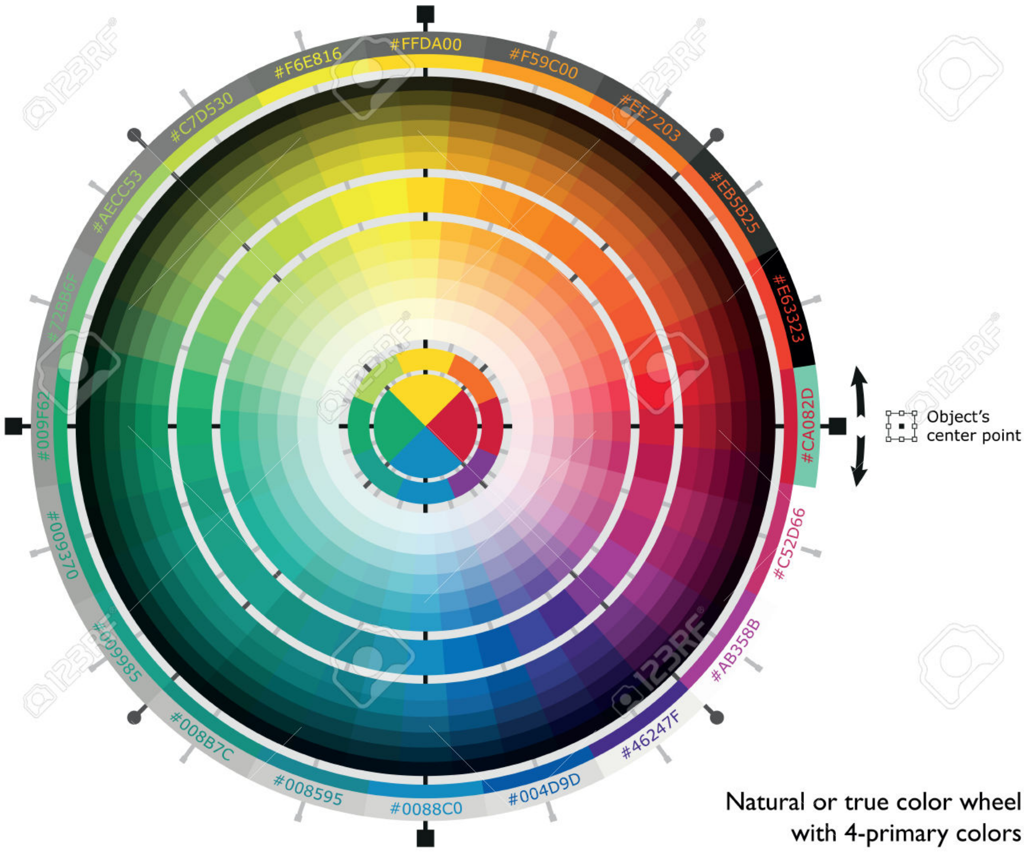

Figure 11: Table 2