SYNAESTHESIA AND LEARNING
A BIDIRECTIONAL RELATIONSHIP

by

MARCUS ROBERT WATSON

B.Hum., Carleton University (2001)
M. Phil., Simon Fraser University (2009)

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

THE FACULTY OF GRADUATE STUDIES

(Psychology)

THE UNIVERSITY OF BRITISH COLUMBIA
(Vancouver)

July 2013

Abstract

I present new evidence about the relationships between learning and *synaesthesia*, particularly *grapheme-colour synaesthesia*, in which individuals experience letters and numbers as coloured. As part of the largest survey of synaesthetic tendencies ever performed, I show that second language acquisition can act as a trigger for the development of synaesthesia, such that children who learn a second language in grade school are three times more likely to develop synaesthesia as native bilinguals. I also demonstrate that previous reports of a sex bias in synaesthesia are almost certainly due to response and compliance biases, rather than any real differences in the prevalence of synaesthesia between men and women. In a detailed examination of the influences of learning on synaesthetic experiences, I show that synaesthetic colours are influenced by knowledge about letters' shapes, frequencies, alphabetical order, phonology, and categorical qualities. Finally, I demonstrate that synaesthesia can itself be exploited in learning. All these results are presented as supporting a *developmental learning hypothesis* of synaesthesia, in which synaesthesia develops, at least in part, because it is useful.

Preface

This thesis describes a number of studies of grapheme-colour synaesthesia that took place at the University of British Columbia, Simon Fraser University, and Charles University in Prague. For the SFU/CU synaesthesia survey, which is the subject of Chapter 2 and provided the data for Chapters 4 and 5, I participated in the research design from the start, in roughly equal collaboration with Kathleen Akins and Lyle Crawford. I did little data collection, but performed all analyses solo, with helpful suggestions and input from collaborators. All the writing in these chapters is my own.

Chapter 3 is a slightly adapted version of a previously published paper, (Watson, M. R., Akins, K. A., & Enns, J. T. (2012). Second-order mappings in grapheme-color synesthesia. *Psychonomic Bulletin and Review*, 19(2), 211-217). The data for this paper was kindly provided by Michael Dixon and Jonathan Carriere of the University of Waterloo. The initial research question was collaboratively arrived at by the three co-authors, all analyses were my own, and I was the principal author of the paper.

Chapter 6 is also taken from a previously published paper (Watson, M. R., Blair, M. R., Kozik, P., Akins, K. A., & Enns, J. T. (2012). Grapheme-color synaesthesia benefits rule-based category learning. *Consciousness and Cognition*, 21, 1533-1540). Here I was the primary person involved in determining the research question and experimental method, though all my collaborators made many useful suggestions and changes. All analysis was performed by myself, and I was the primary author of the paper.

The research described here was approved by the UBC and SFU Offices of Research Ethics. (UBC BREB # H10-00287, SFU # 39456).

Table of Contents

List of Tables

List of Figures

Acknowledgements

I suppose it is standard to say that I couldn't have completed this thesis without the work of many others, but it happens to be true. First of all, my PhD. supervisor Jim Enns did not just improve the specific work I present here, but he taught me how to be a better and smarter researcher overall. His clarity of vision and ability to mull over every aspect of an experimental design or a data set have repeatedly shown me how to get better results and get more out of the results I do have. The five years I have spent working with him have given me the opportunity to see first hand what a scientist looks like, and I only hope to be able to live up to his standard in the rest of my career. He's also a really good guy.

My M.A. supervisor Kathleen Akins has also been a central part of every project in these pages, including being the P.I. on the work in Chapters 2, 4 and 5. She has an astonishing ability to create a theoretical looking-glass world, where what were previously outlandish speculations begin to look normal, and what was previously sober and careful thinking begins to look like sheer madness. I find myself here defending parts of a theory that she first described to me six or seven years ago in her office, which initially seemed bizarre in the extreme but now seems to be one of the most plausible ways of accounting for half the findings in our field. Her friendship and support well above and beyond the call of duty will always be remembered.

Mark Blair has also given me far more support of every kind than I ever deserved. He only directly collaborated on Chapter 6 in this thesis, but he has been a friend, mentor, collaborator and Starcraft opponent for many years now.

Jan Chromý single-handedly took charge of the data collection in the Czech Republic for Chapters 2 and 5, and has been incredibly patient with his substantially less-organized collaborators in Vancouver. We had a far larger team at SFU, and so I owe many thanks to Lyle Crawford, David Gaber, Nazim Keven, Jason Leardi, Lydia Du Bois, Chris Spiker, Nicole Pernat, and others. Lyle Crawford was also a incredibly useful co-author on the paper that outlined the theoretical motivation for this thesis (Watson, Akins, & Crawford, 2010).

John Alderete developed the phonological similarity measures used in Chapter 5, and very patiently explained to a linguistic ignoramus like myself how to think of them. He and Martin Hahn were excellent sources of information on the Czech language in general, and helped design the appropriate analyses. David Eagleman kindly provided several hours of his time to discuss theoretical issues with me, and re-worked his Synesthesia Battery to enable our Czech study to go forward.

My thesis examination committee also included Janet Werker, Peter Graf, Ara Norenzayan and Karon MacLean. Their questions, comments and suggestions have been tremendously helpful in making this thesis more readable and accurate.

Several undergraduate RAs at UBC helped with these studies as well, in particular Chapter 6: Pavel Kozik (who was also a co-author of this chapter), Sean Povill, Sarah MacDonald, Chris Yeh, Spencer Murch, K. L. Ta, and Igal Sterin all put in many thankless hours running subjects and dealing with my various screw-ups. Finally, the entire UBC Vision Lab, past and present, deserves a great deal of thanks for repeatedly listening to me half-coherently outline my latest results and future plans, and then asking me the questions I needed to hear to improve my thinking.

Funding for this research came from a SSHRC Doctoral Fellowship to myself, a James McDonnell Fellowship Centennial Award to Kathleen Akins, and NSERC Discovery Grants to Jim Enns and Mark Blair.

Dedication

This work is dedicated to my wife Cho. Without her love and support I would never have been able to survive almost ten years of graduate school, and our children William and Julia would not have developed into the happy people they are today. Here's to the next decade, and the ones after that.

Love,

M

1 Introduction

> I cannot express it better than to say that a colored idea appears to [me]. [...] Particularly those things which form a simple series; e.g. *numbers, the days of the week, the time periods of history and of human life, the letters of the alphabet, intervals of the musical scale*, and other such similar things, adopt these colors.
>
> These introduce themselves to the mind as if *a series of visible objects in dark space, formless and noticeably of different colors*.
>
> [...]
>
> In the *alphabet*, *A* and *E* are vermillion, *A* however is more cinnabar, *E* is more inclined to rose; *I* is white; *O* orange; *U* black; *Ue* (ü) gray; *C* pale-ash-colored; *D* yellow; *F* dark gray; *H* is bluish ash-colored; *K* nearly dark green (uncertain); *M* and *N* white; *S* dark-blue; *W* brown.
>
> (Jewanski, Day, & Ward, 2009, pp. 297-298, italics in original)

This quotation comes from the first published account of *synaesthesia*, from the doctoral dissertation of George Tobias Ludwig Sachs in 1812, but it could easily have been written by any of millions of people alive today. The final paragraph specifically describes his *grapheme-colour synaesthesia*, in which individuals experience letters or numbers as having colours. There are numerous other varieties, including such oddities as swimming styles that have colours (Nikolic, Jurgens, Rothen, Meier, & Mroczko, 2011), words that have tastes (Cytowic, 1993), calendars and number sequences that lie along convoluted three-dimensional paths in one's personal space (Sagiv, Simner, Collins, Butterworth, & Ward, 2006), letters and numbers that have well-defined personalities and genders (Amin et al., 2011), music that has colour and texture (Head, 2006; Ward, Tsakanikos, & Bray, 2006), and even orgasms that have colours (cf. Novich, Cheng, & Eagleman, 2011).

In the century after Sachs' account was published, synaesthesia became a popular topic of research, attracting the attention of several important figures in early psychology (e.g. Binet & Philippe, 1892; Calkins, 1893; Claparède, 1900; Flournoy, 1892; Galton, 1883). It fell out of favour for much of the 20th century, likely because a condition in which individuals describe unusual internal states without obvious behavioural correlates was incoherent according to the dominant behaviourist framework. Publications on synaesthesia slowed to a trickle (with notable exceptions such as Marks, 1975) from the mid-1930's until Richard Cytowic's work in the 1980's (Cytowic, 1988, 1989a, 1989b; Cytowic & Wood, 1982a, 1982b). His push to bring synaesthesia back to the scientific mainstream came at exactly the right time, and other researchers slowly began to pick up the topic. Behaviourism was long gone, there was a renewed interest in conscious states, and researchers could tackle synaesthesia with the new tools of cognitive neuroscience, and with new insights from the study of sensory development and other unusual conditions such as autism or phantom limb syndrome (e.g. Baron-Cohen & Harrison, 1997; Baron-Cohen, Wyke, & Binnie, 1987; Maurer, 1993; Paulesu et al., 1995; Ramachandran & Rogers-Ramachandran, 1996). Since the turn of the millennium, the floodgates have truly opened, with research groups all over the world studying every aspect of synaesthesia using the full range of tools and methodologies of modern cognitive science, producing almost 500 publications over the past ten years.

My collaborators and I have been a small part of this flood of new research for the past six years, and this thesis brings together our work. I have been particularly interested in the bidirectional relationship between synaesthesia and learning: how learning is a necessary part of the development of synaesthesia, and how synaesthesia might itself be useful for learning. This interest arose because both directions of this relationship are necessary parts of a *developmental learning hypothesis* of synaesthesia developed in collaboration with Kathleen Akins and Lyle Crawford, which states that synaesthesia develops, at least in part, as a strategic aid to overcoming a number of learning challenges in childhood (Watson, Akins, & Crawford, 2010). We initially thought of this as an entirely new idea, but later discovered that important aspects of it had been sketched over a century ago (Calkins, 1893).

This thesis presents evidence for three critical components of the developmental learning hypothesis:

1. Synaesthesia typically develops as part of a difficult learning process in which the synaesthete learns a category structure whose members become the triggers of synaesthetic experiences.
2. Synaesthetic experiences are shaped by this learning process, such that they encode a wide range of information about the learned domain.
3. Synaesthesia is exploited on a variety of memory, learning, and creative tasks, leading to "synaesthetic styles" of performance on these tasks.

The developmental learning hypothesis connects these three claims in a causal chain. Synaesthesia develops in response to various learning challenges (#1) *because* it is useful for these challenges (#3). Synaesthetic experiences are shaped by various aspects of the learned domain (#2) because they developed as part of a learning strategy to acquire knowledge of this domain (#1 and #3).

Establishing this causal chain would prove the developmental learning hypothesis, but is well beyond the scope of this thesis, as it would require a comprehensive set of developmental studies. Rather, my collaborators and I have provided new evidence for each of the three claims, which is presented in the research chapters of this thesis.

In the remainder of this Introduction I want to give the reader the background and context necessary to judge the work I present in the research chapters. I begin with an overview of the current debates over how to define and operationalize synaesthesia, and then turn to what other research has already established about the three claims about synaesthesia and learning.

1.1 What is synaesthesia?

1.1.1 Terminology and definitions

Researchers usually use the term *inducer* to refer to the "trigger" of synaesthetic experiences, and the term *concurrent* to refer to these unusual experiences themselves. Thus George Sachs has letters as inducers and colours as concurrents. The different varieties of synaesthesia are typically named using the formula *inducer-concurrent*, and so we speak of *grapheme-colour*, *music-colour*, or *word-taste* synaesthesias (although this formula may be falling out of favour, cf. "coloured sequence synaesthesia" from Novich, Cheng, & Eagleman, 2011). Coloured inducers are often referred to as *photisms*.

The word *synaesthesia* means "union of the senses", and it was generally thought of in these terms until quite recently. Typically, synaesthesia was defined as a case of unusual associations between sensory modalities, so for example you might have the sense of hearing (music) leading to visual experiences (colour). This cross-modal definition was commonplace despite the fact that grapheme-colour synaesthesia, by far the most studied variety, blatantly contradicts it, as both graphemes and colours are visual. More recently it has become commonplace to address this issue by complicating the definition slightly, e.g. stating that the concurrent experiences are "in another modality [or in] a different aspect of the same sensory modality" as the inducer (Asher et al., 2009, p. 279), which, while more accurate, is so general as to be almost useless. (Are there consistent sensory associations which are not synaesthetic, according to this definition?)

Furthermore, neither synaesthetic inducers nor concurrents are necessarily sensory at all. Grapheme-colour synaesthetes, for example, can experience colours corresponding to the answers of mathematical problems, even when these answers are nowhere in the physical stimulus (Dixon, Smilek, Cudahy, & Merikle, 2000; Smilek, Dixon, Cudahy, & Merikle, 2002a). It is also universally acknowledged that colours for graphemes do not typically vary with font or case (although there may be atypical synaesthetes in this respect, cf. Ramachandran & Hubbard, 2001a), suggesting that the inducer is not a simple sensory stimulus (consider, e.g., how little the shapes of *A* and *a* have in common). Several researchers now explicitly reject the notion that there need be any straightforward sensory aspects to synaesthesia at all, arguing that it is generally triggered by “higher-level” conceptual or linguistic constructs (e.g. Jürgens & Nikolic, 2012; Simner, 2012a). And in the case of varieties of synaesthesia such as graphemes with personalities, it is not clear that the concurrent is sensory either.

The original definition is clearly unsatisfactory, then, but nothing has arisen to take its place. Indeed, a recent exchange of papers has a number of the leading researchers in the field agreeing with each other that we do not know how to define synaesthesia (Cohen Kadosh & Terhune, 2012; Eagleman, 2012; Simner, 2012a; Simner, 2012b). I would argue that this is a feature, not a bug, of modern synaesthesia research. Synaesthesia is a relatively rare and still rather poorly-understood phenomenon, and we run far less risk of needless errors if we do not impose artificial restrictions on what does and does not count as “real” synaesthesia, instead letting the data itself slowly shape our category boundaries.

1.1.2 Operationalizing synaesthesia

While there is no accepted conceptual definition of synaesthesia, there are more-or-less widely-accepted criteria for establishing that a given person is or is not synaesthetic for the purposes of research. There are three broad types of such operationalizing criteria: tests of *self-report*, *consistency*, and *automaticity*.

The self-report criterion is the most most basic, indicating merely that participants self-identify as having synaesthetic experiences. That is, when asked questions like those from Appendix 1, e.g. “When you see, hear, or think about certain letters or numbers, do you see or feel any colours?”, people who answer “no” are not typically considered as synaesthetes. Some criterion of self-identification is almost universal in synaesthesia research, in fact it is so standard that it is frequently not explicitly mentioned by researchers. Nevertheless, it is a crucial test that eliminates roughly 80% of the population from consideration as a synaesthete (see Chapter 2 for rates of self-report of synaesthetic experiences).

The second criterion uses high performance on a consistency test of inducer-concurrent associations as a marker of synaesthesia. In these tests, participants are presented with a series of synaesthetic inducers, usually in random order, and are asked to report their concurrents for each one. At some later time, they are asked to perform the same task again, usually in a new order, and a measure of consistency is taken between responses to the two tests. Those individuals who meet a certain threshold of consistency are deemed to be synaesthetic. The specific details vary quite widely. For example, consistency tests for grapheme-colour synaesthesia might ask participants to report the name of the colour they experience for a given letter (Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996), or to choose the best match to their concurrent from a small sample of colours (Simner et al., 2006), or to choose the specific shade of their experience from the >16,000,000 colours available on a standard computer monitor (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007). The test-retest interval varies from a matter of seconds (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007) to many months (Simner et al., 2006). The precise consistency threshold and definition of consistency differ between studies as well. Typically, less than 5% of the population meets this more stringent requirement.

Another common way of verifying synaesthesia is to test for the automaticity of the inducer-concurrent relationship. Typically, such tests employ a modified Stroop paradigm, in which participants are asked to name either the physical colour of a stimulus or their synaesthetic colour for this stimulus, where the physical colour is either congruent or incongruent with the synaesthetic colour (Dixon, Smilek, Cudahy, & Merikle, 2000; Mills, Boteler, & Oliver, 1999). Usually there is a strong benefit of congruency for response time. This is the least commonly-used criterion of synaesthesia, likely because custom-generating the stimuli for each participant is time-consuming.

Many researchers hope for a clear neurobiological operationalization of synaesthesia (e.g. Simner, 2012a), but for the time being none exists, and so we are left with the three criteria of self-report, consistency and automaticity of synaesthetic experiences. It is important to note that these criteria are independent of each other, not just logically but also in practice. Thus a constant irritation for synaesthesia researchers is that the large majority of individuals who report synaesthetic experiences do not meet the consistency criterion (see Chapter 2). Further, the automaticity criterion can be met without the self-report criterion, as the implicit learning of novel letter-colour associations can lead to performance differences on Stroop-type tasks without any conscious awareness of these associations (Colizoli, Murre, Rouw, Karniel, & Witthoft, 2012; Kusnir & Thut, 2012). After more extensive associative training over the course of weeks, participants may report experiences that sound somewhat similar to conscious synaesthetic concurrents, and this is associated with a stronger Stroop effect, but this has not been systematically explored beyond one study that did not directly address synaesthesia (MacLeod & Dunbar, 1988).

There are theoretical and practical limitations to these criteria, then, but they are the only ones we have. The studies reported here, like the majority of recent work on synaesthesia, use the self-report and consistency criteria: anyone who says they have synaesthetic experiences and is able to consistently reproduce highly similar colours for the same stimuli is treated as synaesthetic. Details of the specific questions asked and consistency test used are found in Chapter 2 and Appendix 1.

1.1.3 Sub-types of synaesthesia

There are varieties of synaesthesia that co-occur in individuals more than others, leading Novich, Cheng, and Eagleman (2011) to identify five distinct sub-types of synaesthesia that cluster in this way. For example, any type of what Novich et al. refer to as *coloured sequence synaesthesia*—grapheme-colour, weekday-colour, etc—is much more likely to co-occur in one individual with another form of coloured sequence synaesthesia, but is not nearly as likely to co-occur with other clusters of synaesthesias, which include *coloured sensations* (e.g. touch-colour or orgasm-colour), *spatial sequences* (e.g. calendar or number forms in personal space), *coloured music*, and synaesthesias with *non-visual sequelae* (e.g. word-taste or sound-smell). Nevertheless, individuals with coloured sequence synaesthesias are far more likely to have, e.g., spatial sequences than individuals with no other forms of synaesthesia (Brang & Ramachandran, 2011; Sagiv, Simner, Collins, Butterworth, & Ward, 2006, see also Chapter 2). So synaesthesias of widely different types appear to be somewhat related, but there are at least five distinct sets of synaesthesias that are especially tightly linked together, for as-yet unknown reasons. This places a potentially important limit on the results in the research chapters of this thesis, as virtually all the analyses reported here deal with some form of coloured sequence synaesthesia.

Researchers have also tried to sub-divide synaesthesia based on differences in phenomenological reports. In particular, there is a widely used distinction between *projector* and *associator* synaesthetes , where projectors are supposed to experience their concurrents as spatially located outside the body, while associators are supposed to experience them without a location in particular. Several researchers have shown that these phenomenological reports correlate with differences in performance (e.g. Dixon, Smilek, & Merikle, 2004) and neurophysiology (e.g. Rouw & Scholte, 2010). However other researchers, including me, have had substantially more difficulty in establishing which group their participants should be classified in. I am confused by several of the questions on a supposedly clear questionnaire to distinguish between projectors and associators (Skelton, Ludwig, & Mohr, 2009), and after using it with a number of participants I found that they had the same confusion, and that responses were not consistent across repeated presentations of the questionnaire. Thus I gave up attempting to classify my participants in this manner. This is not to imply that the classification has no merit, simply that there is still work to be done before it is usable by all researchers. There is evidence that the divisions run further than projector/associator, and that we should categorize at least four different phenomenologies of concurrents (Ward, Li, Salih, & Sagiv, 2007), but for the time being I do not differentiate projectors, associators, or any other phenomenological sub-types of synaesthetes.

1.2 Why are people synaesthetic?

If one wants to understand why there are synaesthetes, there are at least two intertwined questions that are really being asked. The more proximal question asks what it is about synaesthetes' brains that causes inducers to give rise to synaesthetic concurrents: what is the neurophysiology of synaesthesia? The more distal question asks how it is that such brains came about: how does synaesthesia develop?

1.2.1 The synaesthetic brain

All major theories of the neurophysiology of synaesthesia agree that inducers and concurrents are represented by activation in distinct populations of neurons in both synaesthetes and non-synaesthetes. What makes synaesthetes unusual, according to these theories, is that inducer-related activation in area A leads to concurrent-related activation in area B, a causal link that does not occur in non-synaesthetes. There are essentially three broad types of theories about how this happens.

First, *cross-activation* theorists propose that inducer areas have unusually strong direct connections to concurrent areas among synaesthetes, allowing inducer activations to trigger concurrent activations. Thus synaesthesia stems from a breakdown, or at least a reduction, in neural modularity (Baron-Cohen, Harrison, Goldstein, & Wyke, 1993; Maurer, 1993; Ramachandran & Hubbard, 2001b). These theorists have particularly concentrated on grapheme-colour synaesthesia, noting that the so-called *visual word form area* (VWFA) in the fusiform gyrus is right next to the so-called "colour area" V4, which would mean that the unusual connectivity in synaesthesia could be highly localized.

Re-entrant processing theories propose a two-step connection, starting with an area that processes sensory features of the inducer, which projects to an area that processes the concepts or meanings associated with the inducer, which in turn back-projects to sensory areas that represent the concurrent (Smilek, Dixon, Cudahy, & Merikle, 2001). Thus a critical difference between re-entrant and cross-activation models is that in the former the synaesthetic experience is determined by the meaning of the inducer rather than its sensory features. Once again, the model has been most thoroughly fleshed-out with regards to grapheme-colour synaesthesia. As with the cross-activation model, the inducer and concurrent areas in this case are presumed to be VWFA and V4, respectively, while the posterior-inferior-temporal region is suggested as the area representing graphemes' meanings (Smilek, Dixon, Cudahy, & Merikle, 2001).

Finally, *disinhibited feedback* theories propose that the connections between inducer and concurrent areas are no different between synaesthetes and non-synaesthetes, but that there are differences in the feedback from other areas that modulate the signals passing from inducer to concurrent areas (Grossenbacher, 1997; Grossenbacher & Lovelace, 2001). This could either take the form of less inhibitory feedback or more excitatory feedback, in either case making it possible for signals to pass from inducer to concurrent areas that would otherwise be too weak to do so.

Thus far none of these models has been confirmed by actual neurophysiological studies. As Hubbard, Brang, and Ramachandran (2011) point out, the cross-activation theory has probably received the most direct empirical support, but this is patchy at best, and the re-entrant processing and disinhibited feedback models have not been as carefully investigated. This is not to say that there is a lack of neurophysiological studies, just that their results do not map neatly on to any of the theories.

There have now been a number of functional neuroimaging studies showing increased activation of various brain areas during synaesthetic experiences (Aleman, Rutten, Sitskoorn, Dautzenberg, & Ramsey, 2001; Brang, Hubbard, Coulson, Huang, & Ramachandran, 2010; Cohen Kadosh, Kadosh, & Henik, 2007; Hubbard, Arman, Ramachandran, & Boynton, 2005; Laeng, Hugdahl, & Specht, 2011; Nunn et al., 2002; Paulesu et al., 1995; Rich et al., 2006; Rouw & Scholte, 2010; Weiss, Zilles, & Fink, 2005), structural studies showing increased connectivity in various brain areas among synaesthetes (e.g. Banissy et al., 2012; Hänggi, Beeli, Oechslin, & Jancke, 2008; Rouw & Scholte, 2007; Weiss &

Fink, 2009), and EEG/MEG studies showing differences in the time course of neural activity associated with synaesthesia (e.g. Barnett et al., 2008b; Brang, Hubbard, Coulson, Huang, & Ramachandran, 2010; Jäncke, Rogenmoser, Meyer, & Elmer, 2012). (For a review of the brain areas associated with synaesthesia, see Rouw, Scholte, & Colizoli, 2011.)

There are two main ways in which these data fail to map easily on to theoretical predictions. First, results are highly heterogeneous. While there has been some overlap, many differences exist between the various studies, for instance several studies have found that when synaesthetes with coloured concurrents are presented with their inducers they show activation in V4 (e.g. Brang, Hubbard, Coulson, Huang, & Ramachandran, 2010; Hubbard, Arman, Ramachandran, & Boynton, 2005; Nunn et al., 2002) but others have found no such activation (e.g. Rich et al., 2006; Paulesu et al., 1995; Rouw & Scholte, 2010; Weiss, Zilles, & Fink, 2005). Second, many of the areas that do seem to be reliably associated with synaesthesia, such as the precentral gyrus (Laeng, Hugdahl, & Specht, 2011; Nunn et al., 2002; Paulesu et al., 1995; Rouw & Scholte, 2010; Weiss, Zilles, & Fink, 2005), and frontal-parietal networks (Laeng, Hugdahl, & Specht, 2011; Rouw & Scholte, 2010), are not predicted by any of the three models.

The heterogeneity of results is to be expected, given both the wide variety of experimental tasks and tools used and the frequently small sample sizes employed. The heterogeneity of synaesthesia itself is also a serious issue: these studies employ different types of synaesthetes, both in terms of their inducers and concurrents (e.g. grapheme-colour vs. music-colour), and in terms of their self-reported phenomenology (some studies separate associators from projectors, others do not). Until a larger number of experiments are run using consistent methods, it will be hard to conduct appropriate meta-analyses that allow for these inconsistencies to be sorted out.

The activation of areas that are not predicted by any neurophysiological theory of synaesthesia is also not terribly surprising, and does not indicate that any of these theories are false. Rather, it merely shows that they are incomplete.

The developmental learning hypothesis is neutral with regards to the neurological underpinnings of synaesthesia. However it is certainly compatible with an important role for attention and executive functions, as hinted at by the involvement of frontal-parietal networks in synaesthetic experiences (Laeng, Hugdahl, & Specht, 2011; Rouw & Scholte, 2010).

1.2.2 The synaesthetic genome

Whatever the specifics of synaesthetic neurophysiology, how do synaesthetes' brains get to be that way? A genetic explanation for synaesthetic development has been an attractive idea for many researchers. The developmental learning hypothesis is not in conflict with such explanations *per se*. However it would be hard to reconcile it with a *simple* genetic cause of synaesthesia, in which a single gene or group of genes reliably causes synaesthesia to develop (such as is the case for, e.g., Huntington's chorea), since, the developmental learning hypothesis includes an important role for learning in synaesthetic development.

Such simple genetic explanations have been relatively popular, however. There are probably two main reasons for this, one arising from theory and one from evidence. First, a number of researchers have proposed a relatively simple genetic mechanism underlying the high degree of connectivity that is the basis of the cross-activation and re-entrant feedback theories. They suggest that this connectivity may be present in all of us at birth, but deteriorates in the standard process of neural pruning. Synaesthetes, on the other hand, may have a mutation in a gene that controls neural pruning, preventing it from occurring to the same extent as in non-synaesthetes, leading to the unusual connectivity (Baron-Cohen, 1996; Maurer, 1993; Ramachandran & Hubbard, 2001b). Of course this account requires an explanation of why the gene is only selectively expressed in particular regions of cortex, which may complicate the genetic story somewhat.

The more evidence-based reason why researchers began speculating about simple genetic causes for synaesthesia comes from the results of familial studies. Since the earliest days of synaesthesia research, scientists have noted that it runs in families (Galton, 1883) and appeared to be strongly linked to sex. One well-cited study found a female:male ratio of 6:1, further finding that these synaesthetes had a ratio of female:male family members (synaesthetic or not) of 8:1 (Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996). Furthermore, almost all reports of synaesthesia within families involve the trait being passed along the maternal line (Barnett et al., 2008a; Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996; Ward & Simner, 2005). Such skewed ratios require explanation, and a popular hypothesis was that synaesthesia might be an x-linked dominant trait with lethality in males (Bailey & Johnson, 1997; Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996), which might explain both the female bias among synaesthetes and the lack of male family members of synaesthetes.

Two studies effectively ended speculation about the lethality in males of any putative "synaesthesia gene", using much larger samples of families, and finding no difference in the number of male and female family members of synaesthetes (Barnett et al., 2008a; Ward & Simner, 2005). Both studies still found a larger number of female than male synaesthetes, in one case a ratio of 6:1 (Barnett et al., 2008a), and in the other a ratio of 2:1 (Ward & Simner, 2005). The 2:1 ratio was smaller than previous estimates, and there was evidence that even this lower ratio was likely too high due to systematic under-reporting of synaesthesia by men. A later, better-controlled and larger, study (Simner et al., 2006) found no evidence of a female bias at all, and it was argued that the female bias in previous results was largely, if not solely due to differences in response and compliance biases between the sexes (see Chapter 2 for a more complete account of this issue). However it has never yet been firmly established that these response and compliance biases exist.

It seems unlikely, then, that there is a simple x-linked genetic cause of synaesthesia, however it is clear that it does run in families (Barnett et al., 2008a; Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996; Ward & Simner, 2005), suggesting that a genetic component of some kind is at play. Direct comparisons of DNA between synaesthetes and non-synaesthetes have found several candidate chromosonal regions (Asher et al., 2009; Tomson et al., 2011), but these differ between studies, and the same genetic factors are not present in all synaesthetes within either study, suggesting that the genetic influence on synaesthesia is highly polygenic and variable.

1.2.3 The crucial role of learning in synaesthetic development

Most common synaesthetic inducers are culturally transmitted by processes that involve considerable time and effort on the part of the learner, and often formal instruction. In the first paragraph of this thesis I mentioned swimming styles, words, calendars, letters and numbers, music, and orgasms, all but the last of which is clearly learned (and even there one might debate the point). This bias towards learned inducers is not coincidental, and is found in more formal analyses. For example, of the five sub-groups of synaesthesia identified by Novich and colleagues (2011) that were described in the previous section, three of them, which constitute the large majority of cases, exclusively involve learned inducers. Day (2005) lists several dozen types of synaesthesia, including many examples where the inducers are not obviously learned, and certainly not learned deliberately or via formal instruction (e.g. orgasms, smells/tastes, personalities, environmental sounds, temperature). However five of the six most prevalent types of synaesthesia on Day's list involve learned inducers, and these account for the vast majority of his cases. Large-scale surveys of synaesthesia (Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2006 see also Chapter 2 of this thesis) also show an overwhelming majority of cases involving learned inducers.

A simple conclusion follows from this: *synaesthesia normally only develops as part of a formal learning process*. The need to explain this places a serious constraint upon genetic and neurological theories of synaesthetic development. I agree with Cytowic and Eagleman (2009) that so far none of these theories even attempts to do so, because researchers have not generally acknowledged this crucial role of learning, at least not to the the extent that it actually informs their theories (for notable and welcome exceptions to this trend, see, e.g. Simner, Harrold, Creed, Monro, & Foulkes, 2009; Witthoft & Winawer, 2013).

There is only one published account that directly studies the development of synaesthesia in children (Simner, Harrold, Creed, Monro, & Foulkes, 2009). Here a large number (N = 615) of children ages 6-7 were given a modified version of a letter-colour consistency test that required them to select a colour for each letter twice, and then a year later were given the same test. Synaesthetes were identified as those who showed a high degree of consistency in their colours both within each test and across both tests. There was a clear developmental trajectory here: synaesthetes' mean number of consistently-coloured letters was approximately 11 on the first test and 16 on the second.

These results demonstrate that synaesthetic associations coalesce over a lengthy period time that roughly coincides with the development of reading and writing. At 6 years old the average grapheme-colour synaesthete has consistent colours for less than half the letters in the alphabet. One year later, this rises to slightly over half. Clearly, these children have a long way to go before they reach the consistency levels of adult synaesthetes, who frequently have 100% consistent colours.

Like Simner *et al*. (2009), Cytowic and Eagleman (2009, Table 2.2) suggest that the development of synaesthesia coincides with the development of literacy, but they focus on an earlier stage, namely when children first start learning their letters (generally from 34-48 months). If grapheme-colour synaesthesia begins to develop with the first acquisition of letters, then given Simner *et al.*'s (2009) results, the development of grapheme-colour synaesthesia is slow indeed, taking shape over the course of at least six years, and likely much longer. As we will see in the next section, there is evidence that this is exactly what happens, with different stages of learning about letters affecting the development of synaesthesia, leaving behind traces in the synaesthetic colours themselves.

1.3 How does learning change synaesthesia?

We have just seen that synaesthesia generally only develops as part of an explicit learning process, and that this development takes years. What do we know of the influences on this development?

1.3.1 Synaesthetic concurrents as fossils of learning processes

In terms of direct observation, virtually nothing. Simner and colleagues are following the progress of the synaesthetes identified in their childhood study (Simner, Harrold, Creed, Monro, & Foulkes, 2009), but as of yet no further data have been published. However there are a number of papers that describe colour regularities found across adult synaesthetes, particularly grapheme-colour synaesthetes, and I argue that many of these regularities can be thought of as *perceptual/cognitive "fossils"*: traces laid down as a result of the influences on synaesthetic development. Like real fossils, these can tell us a great deal about the environments in which they were formed.

1.3.2 Semantic influences

One class of these involves cases where inducers and concurrents have a common semantic content. For example, *G* is often green for English grapheme-colour synaesthetes, and in general the first letters of common colour words are often associated with the colours named by these words (Barnett et al., 2008a; Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2005; Simner, Harrold, Creed, Monro, & Foulkes, 2009). Similarly, *D* is often brown, which may reflect the fact that *D* is often taught to English speakers as the first letter of "dog", an animal that is stereotypically brown (Rich, Bradshaw, & Mattingley, 2005).

Some of these associations may not be learned until quite late in life. For instance one synaesthete reports that after learning the meaning of the word phthalocyanine (a type of blue-green dye), the colour of the letters within it changed from largely purple and pink to blue and green (Curtis, 1998). (This colour change was only within the context of the word itself - one complicating factor that this thesis ignores entirely is that words often have their own colours that are somewhat independent of the colours of the letters making them up.)

1.3.3 Common associative influences

Another influence on the development of synaesthesia is standard associative learning: some inducer-concurrent associations are formed as a result of the synaesthete being exposed to these associations in the environment. The letter and number colours of some grapheme-colour synaesthetes are derived from those found on toys they played with as young children (Hancock, 2006; Witthoft & Winawer, 2006; Witthoft & Winawer, 2013), others report that they are identical to those used on the wall of their kindergarten (Colizoli, Murre, Rouw, Karniel, & Witthoft, 2012), and other childhood associations have been noted for many years (Calkins, 1893). It should be noted that such easily-determinable associations are relatively rare in the literature, and often attempts to find them result in failure. For example, an in-depth look at the colours of letters in children's books published in Australia between 1862-1989 failed to find evidence of a strong connection to the letter colours of a large sample (N=150) of synaesthetes who grew up there, although there was some evidence that number colours had been influenced by the colours used in a popular method of math teaching during the 1950's and 1960's (Rich, Bradshaw, & Mattingley, 2005).

1.3.4 Universal influences?

There are also common letter-colour associations that may be universal among both synaesthetes and non-synaesthetes. For example, *I* and *O* are most often white for synaesthetes, and X is often black (Simner et al., 2005), and these shape-colour correspondences are found among non-synaesthetes as early as age 2, long before they have started reading (Spector & Maurer, 2008; Spector & Maurer, 2011).

1.3.5 Second-order influences

The semantic, associative, and possibly innate influences on synaesthetic inducers described above are all examples of *first-order mappings*, in which a single element of one domain is related to a single element of another. For example, the individual letter *G* is related to a particular shade of green. These are to be distinguished from *second-order* mappings, “relations between relations”, which, as Chapter 3 explains, are also found in synaesthesia. Here it is not a single element being mapped from one domain to a single element within another domain, but rather a *pattern* or *relationship* within one domain being mapped to a relationship within another domain. For example, synaesthetes have a general tendency to associate similarly-shaped letters, such as *E* and *F*, with similar colours (Brang, Rouw, Ramachandran, & Coulson, 2011; Eagleman, 2010; Jürgens & Nikolic, 2012, see also Chapters 3-5). Here, a relationship of similarity within the domain of shape is mapped on to a relationship of similarity within the domain of colour. These mappings can exist independently of first-order relations - thus *E* might have entirely different colours for different synaesthetes, but if its colour for one individual is similar to *F*'s colour for the same individual, and so on for other synaesthetes, then there is a second-order relationship between shape and synaesthetic colour similarity.

There are also second-order pitch-luminance and pitch hue mappings in music-colour synaesthesia, such that higher pitches tend to be associated with brighter colours (Marks, 1975) and quartertones tend to be associated with colours that are closer to the midpoint of the two adjacent semitones (Head, 2006). Letters and numbers that are more frequently seen in print tend to be associated with brighter colours (Beeli, Esslen, & Jäncke, 2007; Cohen Kadosh, Henik, & Walsh, 2007; Simner & Ward, 2008; Smilek, Carriere, Dixon, & Merikle, 2007), more saturated colours (Beeli, Esslen, & Jäncke, 2007), and colours whose names are more commonly used (Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2005). Letters that appear earlier in the alphabet tend to have more distinct colours from each other than do letters that appear later in the alphabet (Eagleman, 2010, see also Chapter 3). Finally, a number of new second-order findings are detailed in Chapters 3-5.

Synaesthetic concurrents, then, contain numerous traces of the factors that influenced their development. To put it another way, synaesthetic concurrents *encode* a wide variety of information about their inducers, both first-order information about specific inducers and second-order information about the relationships between these inducers. Some of this information is learned right at the start of literacy development (e.g. the shape relationships between letters), some of it quite a bit later (the most extreme example presented here being the colour change due to learning the meaning of phthalocyanine, but Chapter 5 will provide further examples).

1.4 Is synaesthesia good for anything?

In popular culture, synaesthesia is often portrayed as a superpower, sometimes quite literally, (Anonymous, n.d.). It seems intuitive to many people that the kinds of experiences synaesthetes describe ought to be useful in some way. Speaking somewhat more formally, synaesthetes' experiences of their inducers are different from non-synaesthetes, in that they have additional associations with these inducers that the rest of us do not. Do these additional experiences provide any benefit to the synaesthete? Researchers have suggested that this is the case for well over a century (Calkins, 1893), but very little controlled research was conducted into the utility of synaesthesia until the turn of the millennium.

1.4.1 Synaesthesia and memory

Anecdotal reports of synaesthesia's utility for memory are common. Synaesthetes often report using their concurrents in everyday life to assist with remembering names, telephone numbers or the spellings of words (Chapter 2 in Cytowic, 2002, provides several examples), and they tend to self-report better than average memories (Yaro & Ward, 2007). Several savants also report that their synaesthesia is an integral part of their astounding recall, allowing them to memorize *pi* to over 20,000 decimal places, perfectly recall a list of random words on a surprise test 20 years after the initial encoding, or memorize several 50-digit matrices in a matter of minutes, retaining them for months (respectively, Bor, Billington, & Baron-Cohen, 2007; Luria, 1968; Smilek, Dixon, Cudahy, & Merikle, 2002b). Research has confirmed that synaesthesia is associated with mildly enhanced memory for certain stimuli, such as lists of words, but not generally with truly exceptional abilities (Rothen, Meier, & Ward, 2012 provide a detailed overview of the work on synaesthesia and memory), and there is some evidence that this enhancement is due to the synaesthetic concurrents themselves - e.g. synaesthetes' memory for letters can be impaired if these letters are physically coloured incongruently with their synaesthetic colour (Radvansky, Gibson, & McNerney, 2011). Thus while synaesthesia may be an integral part of certain exceptional memory abilities, such abilities are not an integral part of synaesthesia.

At present it is unclear how much of the mild memory benefit associated with synaesthesia is due to the synaesthetic experiences themselves. Studies initially found that grapheme-colour synaesthetes have enhanced memory for colours (Yaro & Ward, 2007) and that number-form synaesthetes have enhanced visuospatial memories (Simner, Mayo, & Spiller, 2009), but that these memory advantages were domain-specific such that, e.g., number-form synaesthetes did not have any enhancement for colour memory. The suggestion was that this would give synaesthetes a memory advantage for those stimuli that induce synaesthesia, but not for other stimuli. However the overall evidence for this is murky at best, as memory advantages can be found on stimuli that do not induce synaesthesia, and some stimuli that do induce synaesthesia do not have a corresponding memory benefit (Rothen, Meier, & Ward, 2012).

1.4.2 Synaesthesia and creativity

A large number of artists are self-reported synaesthetes. There are, for instance, detailed accounts of the synaesthesia of the painters Vassily Kandinsky (Ione, 2004) and David Hockney (Cytowic, 2002); the composers Alexander Scriabin (Peacock, 1985) and Oliver Messiaen (Bernard, 1986); the novelist Vladimir Nabokov (Nabokov, 1989); and the pop/rap musicians Pharrell Williams (Seaberg, 2012) and Kanye West (Anonymous, 2011). Many less famous synaesthetes also offer anecdotal reports of unusually strong artistic interest or ability.

This connection between synaesthesia and creativity has some empirical corroboration. The rate of synaesthesia among fine arts students is 3-4 times as high as that in the general population (Rothen & Meier, 2010b), and synaesthetes are much more likely to be employed in artistic professions than non-synaesthetes (Rich, Bradshaw, & Mattingley, 2005; Ward, Thompson-Lake, Ely, & Kaminski, 2008). Furthermore, synaesthetes tend to score unusually high on the Remote Associates Test, a common measure of creativity (Ward, Thompson-Lake, Ely, & Kaminski, 2008), and members of the general population who score highly on this test also tend to have unusually high consistencies on tests of associations between colour and tones, vowel sounds, and emotional words, although no formal test of synaesthesia was used in this study (Dailey, Martindale, & Borkum, 1997).

As with the memory studies, however, care should be taken not to over-interpret these results: synaesthetes do not out-perform controls on all measures of creativity, and their performance is higher than average, but not extraordinarily so. Once again, it is as yet unclear whether synaesthetic experiences themselves are actually exploited in enhanced creativity, or whether synaesthetes are simply more involved in the arts because their experiences of, e.g., unusual colours lead them to be interested in colour for its own sake (Ward, Thompson-Lake, Ely, & Kaminski, 2008).

1.4.3 Looking for synaesthetic styles of performance

There is some evidence, then, that synaesthesia contributes to enhanced memory, artistic abilities, and other skills, albeit usually to a moderate extent. However synaesthesia serving a useful function does not necessarily entail that synaesthetic performance on everyday tasks is superior to that of non-synaesthetes. Rather, it might be that synaesthesia enables *different* methods of learning, creative production, and so forth, but that non-synaesthetic methods may be just as effective. For instance, many synaesthetes report that their synaesthesia contributes to their memory, but others do not, and there is no difference on standard memory tasks between these two groups (Rothen & Meier, 2010a).

The question, then, is not whether synaesthetes are better than non-synaesthetes at a given task, but rather whether synaesthetes perform this task in a unique manner that exploits their synaesthesia. In this case synaesthesia would still be useful, but not necessarily superior.

Anecdotal reports of a "synaesthetic approach" to memory, learning, or creative endeavours are extremely common. For example,

> Specific numbers interact with other numbers in different ways and their personalities meshed in specific ways as well. When it didn't mesh or color correctly, I knew that I had a) done something wrong or b) had seen a new kind of math. For example, a word like 'pottery' has some of its personality from the double 't' and 'y'. If I left off a 't' or 'y', or spelled it with an 'ie', the personality was off. [...] With music, I could easily tell if a note was off on my viola or if a classmate played the wrong note. It was just the personalities of the notes. (J. Alsaied, personal communication, March 18, 2013)

Another compelling anecdote comes from a synaesthete who reports that when she tried to learn the piano, she discovered that she had three different and inconsistent sets of colours for letters (i.e. the names of the notes), for each of her fingers, and for musical pitches, which made it entirely impossible for her to succeed (Pautzke, 2010). More generally, many of the artists noted previously report using their synaesthesia in their artwork (Anonymous, 2011; Bernard, 1986; Cytowic, 2002; Ione, 2004; Pautzke, 2010), and as previously described many individuals report using their synaesthesia to help with memory and learning (cf. Cytowic, 2002).

There is little experimental evidence for different synaesthetic approaches to tasks, but this is unsurprising, as very little research into this area has been undertaken. Some early work provided in-depth descriptions of a synaesthete's unique approach to a variety of memory and learning tasks, but this was using the method of introspection and was never verified in any other way (Wheeler & Cutsforth, 1921, 1922, 1925). More recently, a synaesthetic savant was shown to actually perform worse than controls at remembering a matrix of numbers when they were coloured incongruently with her synaesthetic experiences (Smilek, Dixon, Cudahy, & Merikle, 2002b). One should note that this congruency effect is not universal: it was observed in a group of child synaesthetes (Green & Goswami, 2008) but was not observed with larger groups of synaesthetes whose memory for digits is only slightly above average (Yaro & Ward, 2007) or no higher than non-synaesthetes (Rothen & Meier, 2009). One interesting interpretation of these results is that synaesthesia *can* be exploited for unusual performance, but need not be. This suggests that if synaesthesia is useful, as the developmental learning hypothesis supposes, it may be useful in different ways for different synaesthetes.

1.5 An outline of this thesis

Each of the research chapters of this thesis is devoted to providing further evidence for one of the three points introduced earlier:

1. Synaesthesia typically develops as part of a difficult learning process in which the synaesthete learns a category structure whose members become the triggers of synaesthetic experiences.
2. Synaesthetic experiences are shaped by this learning process, such that they encode a wide range of information about the learned domain.
3. Synaesthesia is exploited on a variety of memory, learning, and creative tasks, leading to "synaesthetic styles" of performance on these tasks.

I take it that point 1, while not always acknowledged, is uncontroversial: virtually all synaesthetic inducers are only acquired through a lengthy process of learning, usually in a classroom. However Chapter 2, which presents the largest survey on synaesthetic tendencies yet performed, and the first covering more than one linguistic environment, demonstrates that learning's role in the development of synaesthesia is far stronger than previously known. I show that second language learners who are bilingual from infancy are three time less likely to develop synaesthesia (of any kind) than those who learn a second language in grade school. Thus the age at which inducers are learned and the manner in which learning takes place can determine whether or not one develops synaesthesia, or at least whether one maintains it into adulthood. Furthermore, since this determination occurs in grade school, this pushes the development of synaesthesia quite a bit later than is generally assumed. This chapter also establishes that the previously-reported gender bias in synaesthesia is almost certainly entirely due to differences in reporting rates and compliance with experimental protocols, rather than to actual differences in synaesthetic experiences, which further weakens one of the original motivations for assuming a strong genetic cause for synaesthesia.

There is fairly extensive evidence for point 2, and we have already reviewed the ways in which synaesthetic concurrents are modified by various learned characteristics of the inducing domain. This evidence, however, has generally been presented in a piecemeal fashion, with most research focussing on one or two ways in which inducer characteristics map on to concurrent characteristics, almost always from a first-order perspective. In Chapters 3-5 I simultaneously examine multiple second-order influences on grapheme-colour synaesthetes' colours, showing how different aspects of these colours are responsive to different relationships between letters, and how these influences are independent of each other. Chapters 3 (previously published as Watson, Akins, & Enns, 2012) and 4 outline these effects among two different groups of English synaesthetes, while Chapter 5 shifts the focus to Czech grapheme-colour synaesthesia, showing that Czech synaesthetic colours are highly responsive to a extremely detailed set of characteristics of their graphemes. The Czech data also shows that quite sophisticated knowledge about these graphemes affects their synaesthetic colours, showing that the influence of learning about letters on the development of synaesthesia must continue well into the primary school years.

As we have seen, there is a great deal of anecdotal evidence for point 3, but less direct support for it in the laboratory. Chapter 6 outlines a categorization learning study (previously published as Watson, Blair, Kozik, Akins, & Enns, 2012) showing not only that synaesthetes perform differently from non-synaesthetes looking at the same stimuli (and far better than these non-synaesthetes), but also that these differences in performance arise because the synaesthetes exploit their synaesthetic colours to succeed, and not because of general group differences such as motivation or memory advantages. This is also the first time that synaesthesia has been shown to be useful for a difficult learning task that involves consciously coordinating multiple pieces of information.

Finally, the Conclusion will attempt to tie all these strands back together, and assess the plausibility of the developmental learning hypothesis of synaesthesia in light of the new evidence I have presented.

2 Childhood learning and rates of synaesthesia—The prevalence of synaesthesia in Czech and English[1]

2.1 Introduction

If synaesthesia develops in response to childhood learning challenges, then some differences in these challenges ought to correspond with differences in the development of synaesthesia. One way this might manifest itself would be if synaesthesia were more (or less) likely to develop in children who face a particular type of learning challenge. In this case, one would expect to find differences in rates of synaesthesia in adults that correspond to differences in these challenges. The present study looks for evidence of exactly this, by measuring rates of synaesthesia among adults and seeing if these rates are associated with differences in the particular learning challenges these adults had faced as children. Of course since synaesthesia is a relatively rare condition, testing this requires a fairly large study. Thus it was that over the course of four years we found ourselves conducting by far the largest study of synaesthetic tendencies yet performed (N =

1. The data presented in this chapter and Chapters 4 and 5 was collected over several years at Simon Fraser University and Charles University in the Czech Republic. I participated in the research design from the start, in roughly equal collaboration with Kathleen Akins and Lyle Crawford. Data collection was coordinated by Jan Chromý at Charles University and by myself, Kathleen Akins, Lyle Crawford and a number of research assistants at SFU. I wrote the chapters myself and performed all analyses solo, with helpful suggestions and input from collaborators.

11,664) at Charles University in Prague, Czech Republic, and Simon Fraser University in Burnaby, British Columbia.

2.1.1 The development of literacy and the development of synaesthesia

The particular learning challenge, or more accurately group of challenges, of interest in this study is that associated with becoming literate. We know that grapheme-colour synaesthesia begins developing prior to age 6, and continues to develop after age 8 (Simner, Harrold, Creed, Monro, & Foulkes, 2009), meaning that much, if not all, of its development overlaps with the development of literacy. We also know that synaesthetic colours are influenced by a wide range of learned properties of letters (Beeli, Esslen, & Jäncke, 2007; Brang, Rouw, Ramachandran, & Coulson, 2011; Eagleman, 2010; Jürgens & Nikolic, 2012; Simner et al., 2005; Simner & Ward, 2008; Smilek, Carriere, Dixon, & Merikle, 2007; see also Chapters 3-5), meaning that the development of grapheme-colour synaesthesia is influenced by the development of literacy. What is unclear is how far this influence goes. Is it merely that specific grapheme-colour associations are changed in the course of becoming literate, as is already known, or is there a deeper influence, such that factors that change the process of becoming literate can change the likelihood of developing grapheme-colour synaesthesia in the first place? This study explores the second possibility.

Differences in learning might affect rates of synaesthesia in a number of ways. We offer two competing hypotheses for how this might occur, the *complexity* and *simplicity hypotheses*. The complexity hypothesis states that synaesthetic associations are more likely to develop when children are faced with more complex or difficult tasks, because they would be most useful when task demands are high. The simplicity task, as the name implies, states that synaesthesia is more likely to develop when faced with a simpler task, simpler either because it is intrinsically easier or because the child understands some aspect of the task that makes it simpler for them. Using synaesthesia as a strategy might simply require cognitive resources that are not available when the learning task is too challenging, or might only be feasible when the child has enough conceptual understanding of the inducer domain. For instance, one might need to have a high degree of metalinguistic awareness, understanding what letters are and how an alphabet works, in order to develop grapheme-colour synaesthesia.

There are three specific factors that we reasoned might reliably indicate differences in the difficulty of becoming literate in ways that might also affect the development of synaesthesia. These include the *orthographic transparency* of the language that one is learning to write (how orderly and simple the relationship between phonemes and written graphemes is), the acquisition of second languages by the children who are learning to write, and any unusual strengths or weaknesses that a child has with reading or writing such as learning to read very early or, conversely, dyslexia. For all three factors, we hypothesized they would affect the rates of grapheme-colour but not other types of synaesthesia, since the letters that induce grapheme-colour synaesthesia are, obviously, a critical component of literacy.

Our interest in orthographic transparency led to the study being run in the Czech Republic and Canada. Czech and English are interesting comparison cases because their alphabets are highly similar, but the relationship between these alphabets and the phonology of the spoken language differs greatly. Czech is highly orthographically transparent—there is almost a one-to-one mapping between letter identities and their corresponding phonemes (see Chapter 5 for a more complete description of Czech orthography). English, on the other hand, is about as orthographically opaque as is possible—each letter can produce several different phonemes, and each phoneme can be produced by many different combinations of letters. Learning phoneme-letter correspondences, a crucial part of learning to read, is a more difficult task in orthographically opaque languages such as English than in orthographically transparent languages (Ellis et al., 2004; Seymour, Aro, & Erskine, 2003). According to the complexity hypothesis, then, there ought to be more English-speakers than Czech-speakers resorting to unusual strategies such as using synaesthetic associations in order to learn the difficult English orthography, and thus to higher rates of grapheme-colour synaesthesia among English speakers, while the simplicity hypothesis predicts the reverse.

It is well-established that second language acquisition has associated benefits that affect the development of literacy. In particular, bilingual children have advantages in executive function and metalinguistic knowledge. Either of these, according to the simplicity hypothesis, could lead to higher rates of synaesthesia among bilinguals (and according to the complexity hypothesis could lead to higher rates of synaesthesia among the monolinguals who do not have these benefits). Thus the survey asked participants to indicate what languages they spoke and when they began learning them, enabling us to verify if the acquisition of second languages affects the prevalence of grapheme-colour synaesthesia.

Finally, reading ability might affect rates of grapheme-colour synaesthesia in at least two different ways. First, according to the complexity hypothesis, someone who found reading particularly difficult might be more likely to develop synaesthesia. Conversely, someone with unusually strong reading abilities might be more likely to develop grapheme-colour synaesthesia according to the simplicity hypothesis. Cytowic and Eagleman (2009) also suggest that early reading might be associated with synaesthesia. We asked all participants if they learned to read unusually early (before kindergarten) and also asked several questions about reading difficulties in childhood or at present.

This is, clearly, a highly exploratory study, and before starting it was apparent that we might very well not find any influence of orthographic transparency, second language acquisition, or reading ability on rates of synaesthesia. However testing for these three effects required a large-scale survey of synaesthetic tendencies, indeed the largest yet performed, and so would produce the highest-quality data on the epidemiology of synaesthesia. In particular, we expected to provide clear answers to two long-standing epidemiological questions: how common is synaesthesia, and are there more female than male synaesthetes?

2.1.2 How common is synaesthesia?

Numerous surveys of synaesthetic tendencies have been made over the past 130 years (Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996; Calkins, 1893; Cytowic, 1993; Cytowic, 1997; Domino, 1989; Galton, 1883; Niccolai, 2012; Ramachandran & Hubbard, 2001a; Rich, Bradshaw, & Mattingley, 2005; Rose, 1909; Rothen & Meier, 2010b; Simner et al., 2006; Simner, Harrold, Creed, Monro, & Foulkes, 2009; Ulich, 1957; Ward & Simner, 2005), but there is little agreement among them. Estimates of the prevalence of grapheme-colour synaesthesia, for example, range from 0.05% (Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996) to 13% (Calkins, 1893).

Simner and colleagues (2006) argue that these differences are due to two key methodological flaws, one or both of which are found in almost all these studies. Some studies are far too liberal, in that they simply ask a large group of people whether they make synaesthetic associations and take them at their word, meaning they use the self-report criterion for synaesthesia and nothing else. Most recent studies avoid this problem by using a more stringent consistency criterion, however most of them do not randomly sample the population, instead relying on self-referral by synaesthetes, canvassing subjects by means such as newspaper advertisements. This makes it likely that their rates are far too conservative, as it is very unlikely that anything other than a small minority of newspaper readers would respond to such a advertisement.

To date there have only been two studies that avoid both of these problems. Simner *et al*. (2006) directly asked a random sample of individuals about their synaesthetic tendencies, thereby avoiding the self-referral problem, and immediately followed this up with rigorous tests of consistency, finding an overall prevalence of approximately 4.4% for all varieties of synaesthesia. Rothen & Meier (2010b) provided a grapheme-colour consistency test to everyone in their sample, finding a prevalence of 2% in the general population and 7% among fine-arts students. However even with samples as large as 500, these studies are somewhat underpowered to establish precise rates: for example the University study of Simner *et al*. (2006) finds that grapheme-colour synaesthesia has an estimated prevalence of 1.8%, but the 95% confidence interval for this estimate ranges from 0.6-3.0%. This is still vastly better than the 0.05-13.0% range established from former conservative and liberal estimates (cf., respectively, Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996; Calkins, 1893). However, given that the upper bound of Simner *et al.*'s confidence interval is 5 times its lower bound, it is clear that anyone interested in comparing rates across groups, as we are in the present study, would be unable to find anything other than exceptionally large effects (such as those found by Rothen & Meier, 2010b).

Our survey was handed out to a random sample of students at Charles University, and the online Synesthesia Battery (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007) provided a rigorous consistency test. Thus, like Simner *et al*. (2006) and Rothen & Meier (2010b), the present study avoids both the liberal and conservative flaws of previous epidemiological studies, and its much larger sample size allows for a higher degree of precision in the results.

2.1.3 Are there more female than male synaesthetes?

Several studies have reported a strong female bias among synaesthetes, with reported relative rates of synaesthesia associated with being female as high as 6 (Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996). Ward and Simner (2005) cast doubt on this by conducting a much larger familial study (85 families compared to 6 in Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996), in which they found a far smaller female:male bias than in previous studies. Their initial group of 85 synaesthetes was entirely composed of self-referred synaesthetes, and was strongly female-biased (4:1). These participants identified a second group of 58 family members who they knew to be synaesthetic as a result of directly speaking to them about the topic, and the ratio of female:male synaesthetes was recalculated using both groups to be 2:1. This halving of the degree of female bias suggests that the initial bias of 4:1 was largely due to women being more likely to self-refer themselves, which means that even the lower ratio of 2:1 is likely too high, for two reasons. First, as Ward and Simner (2005) point out, a female referral bias would mean that male synaesthetes who were the sole synaesthetic members of their families were far less likely to come to the attention of the study, as they could only be members of the first group, since they could not be identified as synaesthetic by other family members, yet their lower rates of self-referral would make them less likely to be part of this first group. They do not discuss a second potential reason why their 2:1 ratio may be too high, namely that a female self-reporting bias might also affect the likelihood of family members reporting synaesthesia *to each other*, not simply to researchers, which could skew the composition of their second group of family members.

A more recent large-scale familial study employing a similar methodology (Barnett et al., 2008a) still found a 6:1 female:male ratio of confirmed synaesthetes, but this ratio is highly doubtful. First, as just discussed, male synaesthetes who are the sole synaesthetic members of their families would probably be less likely to come to the attention of the study. Second, the proportion of family members who were able to be contacted by the researchers differed between genders, with 43% of female relatives uncontacted, compared to 69% of male relatives.

Ward and Simner (2005) conclude that accurately determining the ratio of female:male synaesthetes can only be done using a large-scale study that does not rely on self-referral. They proceeded to carry out such a study (Simner et al., 2006), which found a ratio of female:male synaesthetes of 1.1:1, which was not significant. Thus they concluded that previous reports of female biases were largely, if not entirely, due to a female self-referral bias rather than to real differences in the prevalence of synaesthesia, but did not entirely rule out the possibility of a small female bias. Simner *et al.*'s (2009) later study of the development of synaesthesia in childhood found a female:male ratio of 1.6:1, which was also non-significant due to a small sample size, but suggestive nevertheless.

Our study may be able to shed more light on this question. As sample sizes are much larger than any previous study, we have the potential to find small but significant effects.

Furthermore, the study mixes random sampling and self-referral in a way that may allow for a female bias to be observed in action. In the first phase of the study, a paper survey is handed out to a random sample of students in undergraduate university classes, but in the second phase of the study, participants are invited to register for an online consistency test (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007) and complete it to be confirmed as synaesthetes. Hence differences in willingness to comply with the experimental protocol should have relatively little impact on the first phase of the study, but could potentially affect the second phase. Thus we can not only determine sex biases in rates of self-reported and confirmed synaesthesia, but we can also verify if there is a difference in compliance between men and women during the second phase, which will provide further evidence for the source of any gender differences in rates of synaesthesia.

2.1.4 Outline of the study

The study consisted of two phases. First, a paper survey (see Appendix 1) was given to university students at Charles University (hereafter CU) in the Czech Republic and Simon Fraser University (hereafter SFU) in Canada. This survey included descriptions of a number of synaesthetic experiences and asked respondents to indicate if they had these or similar experiences, as well as a number of questions about factors that we thought might relate to synaesthesia (e.g. second-language acquisition, reading ability, and gender). Respondents who answered positively were then invited to participate in the second phase of the study by taking the online consistency tests at the Synesthesia Battery (www.synesthete.org, Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007).

Analyses consisted of calculating the reported and confirmed rates of the various types of synaesthesia, and determining what other factors were associated with each of these types. In particular, we verified whether there were prevalence differences due to native language (which, given our two populations, corresponds to orthographic transparency), second language acquisition, reading ability, and gender. We also determined the degree to which the various types of synaesthesia tend to cluster together in individuals, by calculating whether having a given type of synaesthesia makes one more or less likely to have another type.

2.2 Methods

2.2.1 Phase I - Paper survey

The paper survey (see Appendix 1) included questions about gender, handedness, languages spoken and the age of acquisition of these languages; six questions asking if participants had specific synaesthetic experiences (e.g. “When you see, hear, or think about certain letters or numbers, do you see or feel any colours? Example: There is something yellow about the letter G."); one open-ended question inviting participants to record any synaesthetic experiences not covered by the other questions; one question about using letters or numbers as characters in childhood stories; and four questions on reading abilities. We also asked participants for their email address in order to invite them to participate in phase 2. The survey was in English at SFU and in Czech at CU.

Surveys were handed out to students in undergraduate university classes at the start of the period. A brief presentation was given describing synaesthesia and explaining that participation was strictly voluntary and was not tied to their course performance. Approximately 10 minutes after the surveys had been handed out, they were collected. It was not possible to keep a precise count of the number of students who chose to complete the survey, but we estimate that well over 95% completed and returned the survey in both universities.

5001 students from CU and 6663 students from SFU returned a completed survey. Of these, 3431 CU (69%) and 6084 SFU (91%) students provided an email address, which was necessary for the second phase of the study. A wide variety of courses were sampled in both universities, primarily from Arts and Social Sciences faculties. 69% of CU and 61% of SFU respondents were female.

2.2.2 Phase II - Synesthesia Battery

2394 CU and 2054 SFU respondents (48% and 31%, respectively, of the total survey respondents) reported that they experienced colored letters, numbers, weekdays, months, or sounds. Of these, 1862 CU and 1881 SFU respondents also provided an email address (78% and 92%, respectively, of those who reported synaesthesia). Each of these individuals was emailed and invited to participate in the verification stage of the study using the online Synesthesia Battery (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007), which includes consistency tests for a wide range of synaesthesias. Those who reported other forms of potential synaesthesia were not contacted, as we did not have easily-available online tests of consistency. In exchange for registering, each Battery participant was entered into a lottery for 10,000 Czech Crowns or 500 Canadian dollars (both approximately 500 USD) with a greater than 1% chance of winning. 355 CU and 302 SFU participants registered for the Battery (thus we ran 4 lotteries in each country).

After registering on the Synesthesia Battery website, participants were presented with a checklist of synaesthesia types, and instructed to select any types that they may have. If they selected a form of synaesthesia for which the Battery includes a consistency test, then they were given this test. These tests are described in detail in Eagleman et al. (2007) but in brief, they require participants to choose a color for each inducer in random order, then repeat this process twice more. The similarity between the three colours assigned to each inducer is calculated as a distance in RGB space, and the average across all inducers is calculated. This mean distance is the participant's consistency score for a test, and each participant will generate one consistency score for each of the tests they complete. Eagleman et al. (2007) suggest that genuine synaesthetes tend to have a consistency score below 1, which is the threshold we adopted.

2.3 Results

2.3.1 Linguistic characteristics of the samples

There were large differences between the linguistic capabilities reported by students at CU and SFU (see Figure 2.1). Broadly speaking, the CU population was almost entirely composed of native Czech or Slovak speakers who learned multiple other European languages in grade school. (Czech and Slovak are closely related, largely mutually intelligible languages that use identical alphabets.) The SFU population was far more heterogeneous, containing roughly equal proportions of respondents who were bilingual from birth, who learned a second language from kindergarten on, and who were monolingual. There are increases in the acquisition of a second language in the SFU population at ages 5, 10, and 14+, which likely correspond, respectively, to large numbers of children entering second-language immersion programs in kindergarten, to the beginning of the British Columbia elementary second-language programs in grade 5 (BC Ministry of Education, 2001), and to second-language learning in high school and beyond. Unlike the CU sample, SFU multilinguals spoke languages from every corner of the globe. We adopted a conservative definition of “native” languages: any language reported as having been acquired after their second birthday was considered non-native.

4727 CU respondents (95%) reported speaking Czech or Slovak as a native language, compared to only 4156 SFU participants (62%) speaking English as a native language. There were also far fewer distinct languages reported by CU (59) than SFU (158) students, with only 10 of the CU languages being non-European, compared to 123 at SFU. CU students tended to be more multilingual, with only 5 students (0%) reporting themselves to be monolingual, compared to 1507 SFU students (23%), and a mean number of languages spoken of 3.5 compared to 2.3. However there were far fewer *native* multilinguals in the CU sample, with only 102 students (2%) who reported speaking at least 2 languages before age 2, compared to 1427 (21%) at SFU.

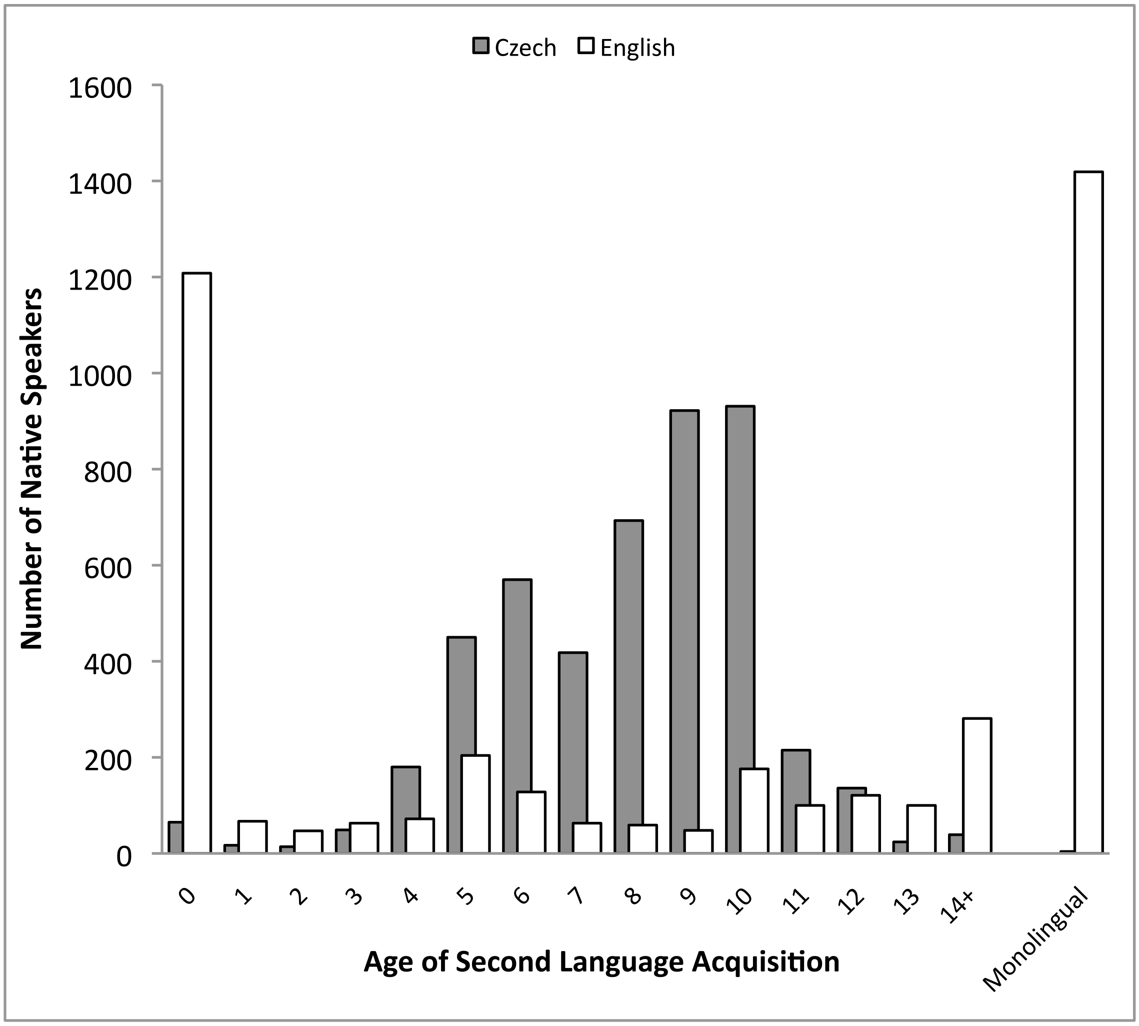

Figure 2.1 Differences in second language acquisition between Czech and English speakers.

Unless otherwise stated, all proportions reported during the remainder of these results are taken from the 4727 native speakers of Czech or Slovak (CU) and the 4156 native speakers of English (SFU).

2.3.2 Higher rate of endorsing synaesthesia among Czechs

Figure 2.2 presents the proportions of native-speaking survey respondents who reported synaesthetic experiences on the paper survey. In general, reported rates were high, with each form of synaesthesia described on the survey being endorsed by 5-40% of respondents at both universities, and 5% reporting other types of experience that they felt might qualify as synaesthetic. The vast majority of responses to the "other types" question appeared to be consistent with those reported in other surveys of the range of synaesthetic experiences (e.g. Day, 2005). Most were associations between particular categories (e.g. school subjects, seasons, tastes, bodily actions, car models, letters, music) and sensory experiences (colors, sounds, tastes, smells, shapes, sizes, textures, and temperatures). There were also many reports of personifications of inanimate objects, as well as of arrangements of various categories in personal space. These reports, however, were not verified any further, and neither were the reports of word-taste, number-form, and grapheme-personality synaesthesia.

All varieties of synaesthesia save the "other" category were endorsed by a greater proportion of CU than SFU respondents, a difference that was significant in all cases save grapheme-personality synaesthesia (as the error bars in Figure 2.2 are 95% C.I.'s, any cases where they do not cross are significantly different). It may be worth noting that differences in the proportions of letter-color and number-color synaesthesia were solely due to differences in the proportion of participants who reported *both* types of synaesthesia. That is, there were no significant differences between the proportion of respondents endorsing letter-color but not number-color synaesthesia (CU: 10.0%, SFU: 8.5%), or the converse (CU: 5.1%, SFU: 4.5%), but there was a large difference between the proportion who reported experiencing both (CU: 6.0%, SFU: 1.5%), and this difference was the main contributor to the overall differences between the two rates.

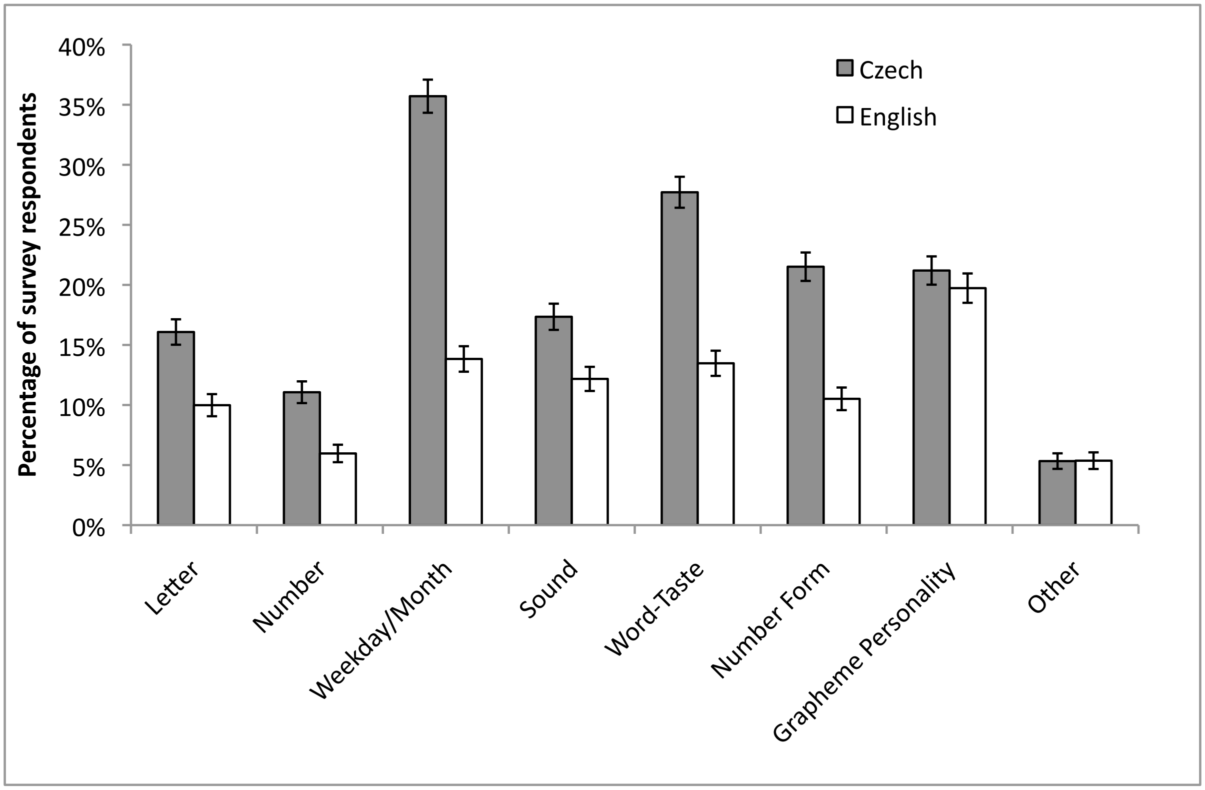

Figure 2.2 Reported rates of synaesthetic experiences for native Czech and English speakers. Error bars indicate +/- 95% C.I.s.

2.3.3 Higher rates of confirmed synaesthesia among Czechs

Figure 2.3 presents the proportions of synaesthesias confirmed by the Synesthesia Battery among native speakers, which ranged from less than 0.05% to 1.8%. The same general trend of higher rates of Czech synaesthesia was observed: absolute proportions were higher among the CU sample than SFU for all varieties of synaesthesia tested on the battery save chord-color synaesthesia, and significantly higher in all cases save letter-colour and scale-colour synaesthesia. Interestingly, the differences in the rates of reported sound-color synaesthesia (see Figure 2.2) appeared to be driven largely by differences in instrument-color synaesthesia. (Since the paper survey asked about "sound-color" synaesthesia in general, and the Battery only includes tests for chords, instruments, and scales, it is possible that there are differences in the rates of synaesthetic colors experienced in response to speech or other sounds, but this could not be verified.) There was a high degree of crossover between the various types of synaesthesia, with most synaesthetes being confirmed as having multiple types (see section 2.3.7 and Appendices 2 and 3 for a complete discussion of this). We found an overall prevalence of 3.7% synaesthetes in the CU sample and 2.3% in the SFU sample.

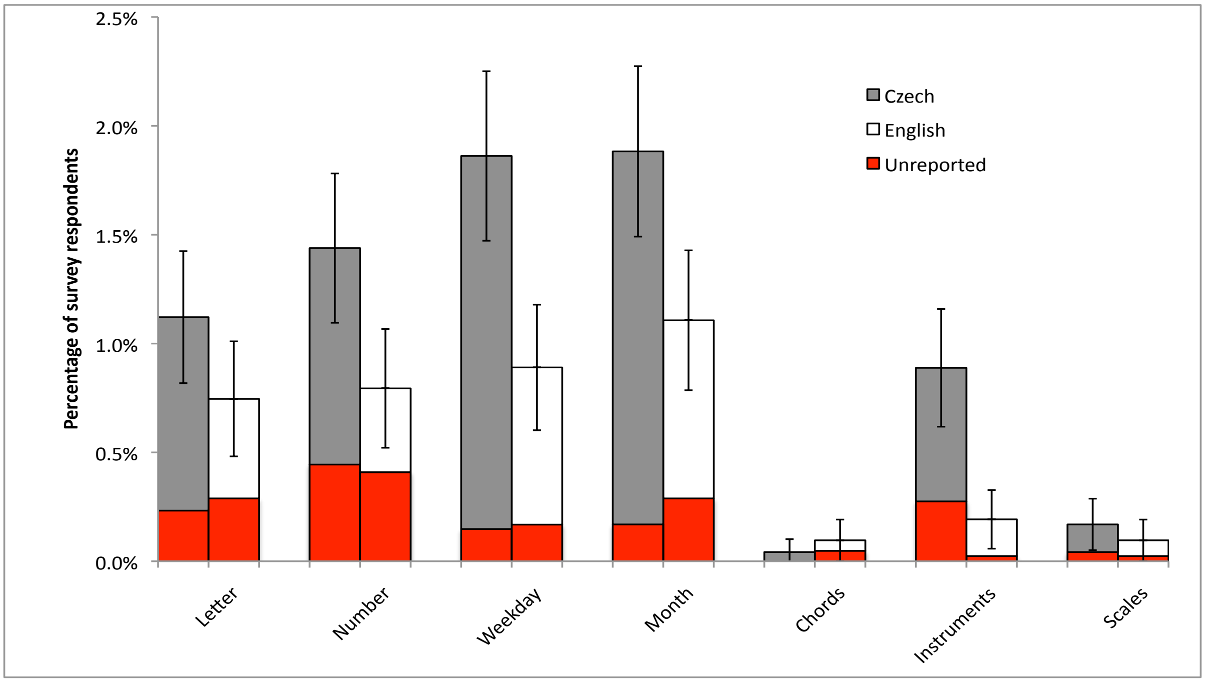

Figure 2.3 Confirmed rates of synaesthesia for native Czech and English speakers. Error bars indicate +/- 95% C.I.s. "Unreported" synaesthetes are those who reported that they did not have the type of synaesthesia in question on the survey, but tested positively for it on the Synesthesia Battery.

An unanticipated finding was that there were a relatively large number of confirmed synaesthetes who reported on the initial paper survey that they did not have the type of synaesthesia in question, as shown in Figure 2.3. For instance, 19.0% (CU) and 30.0% (SFU) of confirmed letter-color synaesthetes selected the response "No, I do not have experiences like this", when asked about associations between letters and colors on the survey. Since only participants who reported synaesthetic experiences were asked to register for the Synesthesia Battery, this means that all of these participants reported having at least one other type of synaesthetic experience. (E.g. 7 of the 8 confirmed CU letter-color synaesthetes who reported not experiencing colored letters did report experiencing weekday or month color synaesthesia, and 3 of them reported experiencing number-color synaesthesia.)

2.3.4 Higher rates among Czechs are due to late second-language acquisition

What explains the greater prevalence of synaesthesia among Czechs? The second factor we suggested might impact rates of synaesthesia is second-language acquisition. Recall that virtually all the Czech speakers learned a second language, but not as a native speaker, whereas English speakers were almost evenly split between monolinguals, native multi-linguals, and non-native multi-linguals (see Figure 2.1). If rates of synaesthesia are higher among non-native multilinguals than among monolinguals or native multilinguals, then this could explain the higher rates among Czech speakers. Figure 2.4 shows the rates of confirmed synaesthesia among these three second-language groups. In order to maximize power, this analysis was performed using all 11,664 survey participants from both samples, regardless of whether they were native speakers or not. Furthermore, since the greater rates among Czechs were found in almost all varieties of synaesthesia we tested, we combined all varieties into a single measure, which indicated if a participant was confirmed as having any one of the varieties of synaesthesia shown in Figure 2.3. Figure 2.4 demonstrates a clear prevalence difference between second-language groups. In particular, those who learn a second language between the ages of 2-12 are 2-4 times as likely to have synaesthesia as native bilinguals, a difference which is significant between ages 5-10. An identical analysis using only SFU participants yielded the same pattern of results. (This could not be meaningfully performed using only CU participants due to the extremely low numbers of Czech monolinguals and native multi-linguals.)

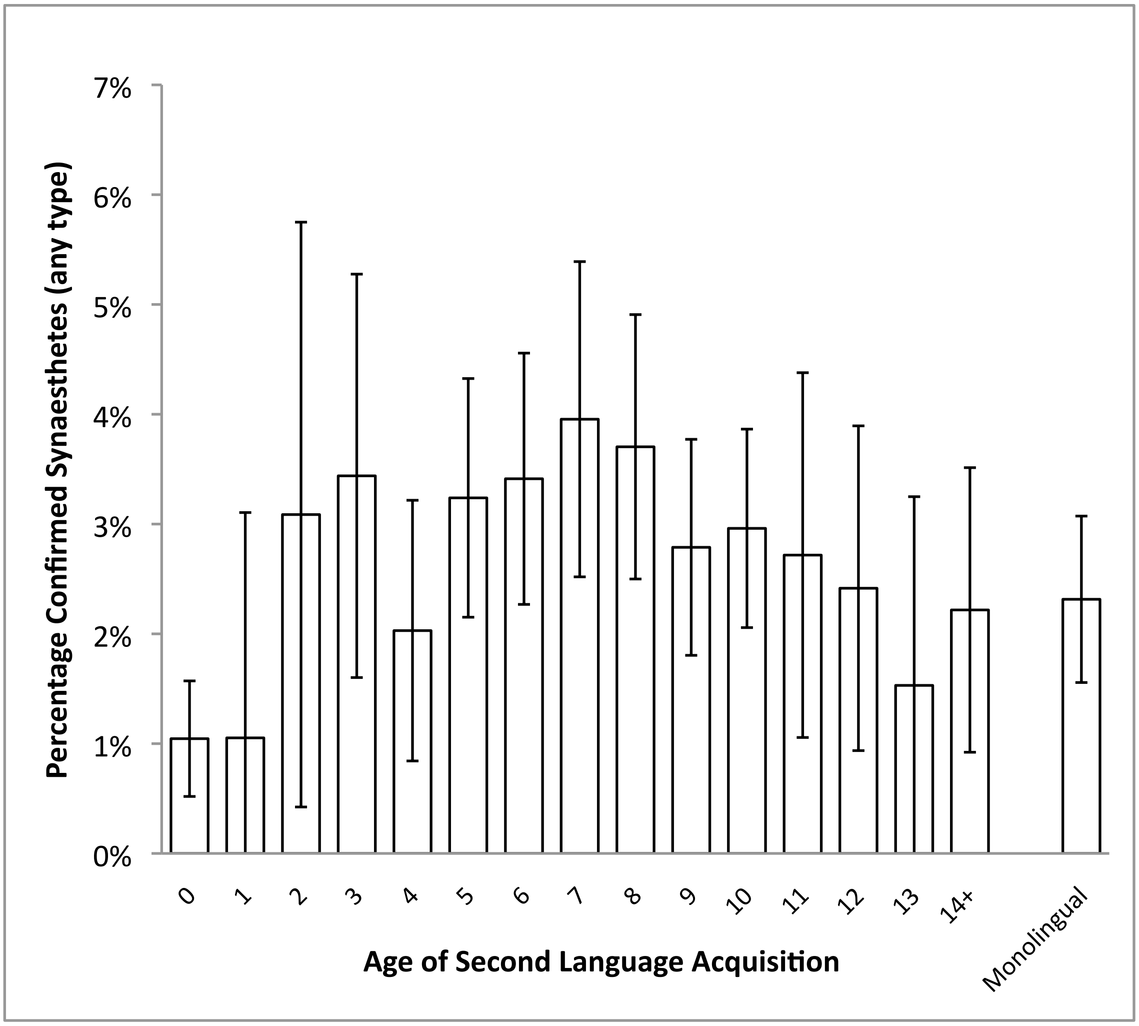

Figure 2.4 Rates of confirmed synaesthesia by age of second language acquisition, for all survey participants in both countries. Error bars indicate +/- 95% C.I.s.

This means than any comparison of all CU participants to all SFU participants, as in Figure 2.3, would confound at least two factors: native language and age of second-language acquisition. Since there were effectively no native multilingual or monolingual Czechs, the only way to test for an effect of native language without this confound was to test only non-native multilinguals (i.e. to eliminate roughly 2/3 of the English speakers from the analysis). When we did this, restricting our sample to native speakers of Czech or English who learned a second language after their second birthday, no significant differences between the two native language groups were found (Figure 2.5).

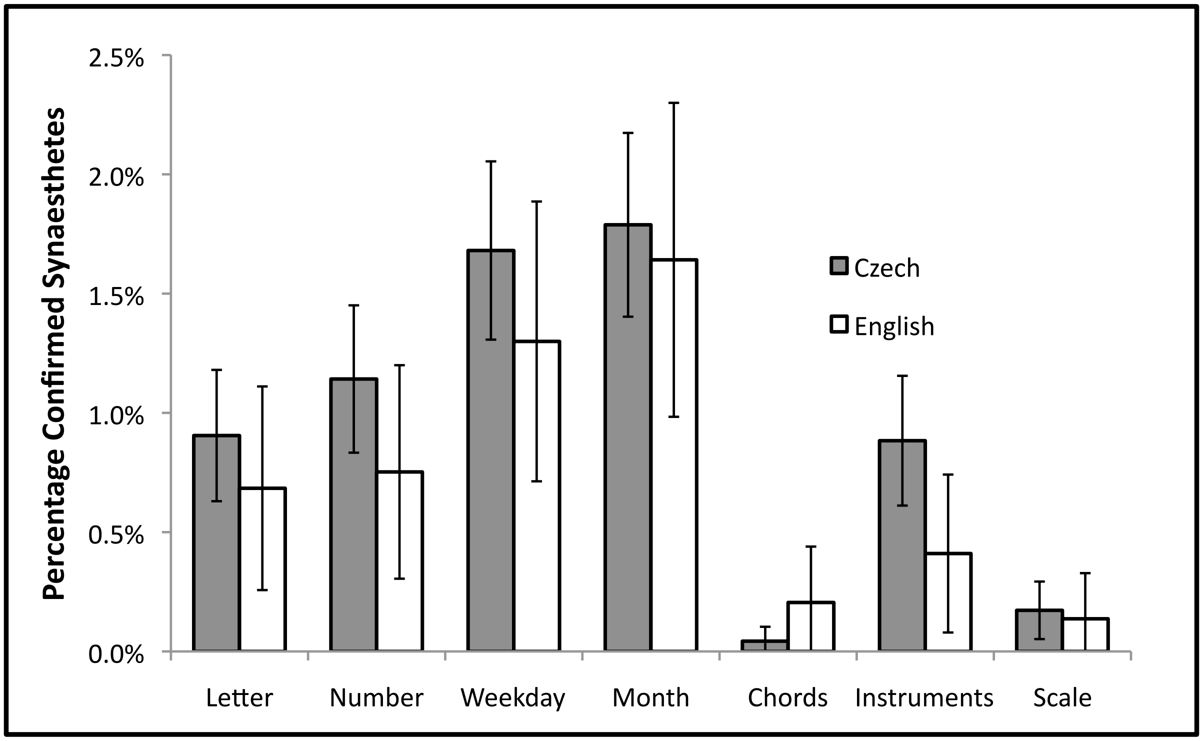

Figure 2.5 Confirmed rates of synaesthesia among native Czech and English participants who acquired a second language after age 1. Error bars indicate +/- 95% C.I.s.

2.3.5 Grapheme stories and reading abilities are associated with reported but not confirmed synaesthesia

Figure 2.6 shows participants' endorsement of the final five statements on the paper survey, which included one statement concerning telling stories involving graphemes, and four concerning reading ability. CU participants were more likely to report telling grapheme stories (24.4%, SFU: 17.3%), while SFU participants were more likely to report learning to read before kindergarten (CU: 43.5%, SFU: 48.8%) and also more likely to report having had special assistance with reading as a child (CU:13.2%, SFU: 20.3%) and having reading difficulties as an adult (CU: 1.9%, SFU: 6.3%). Equal proportions of respondents endorsed having been formally diagnosed with a reading disorder as a child (CU: 6.6%, SFU: 6.2%).

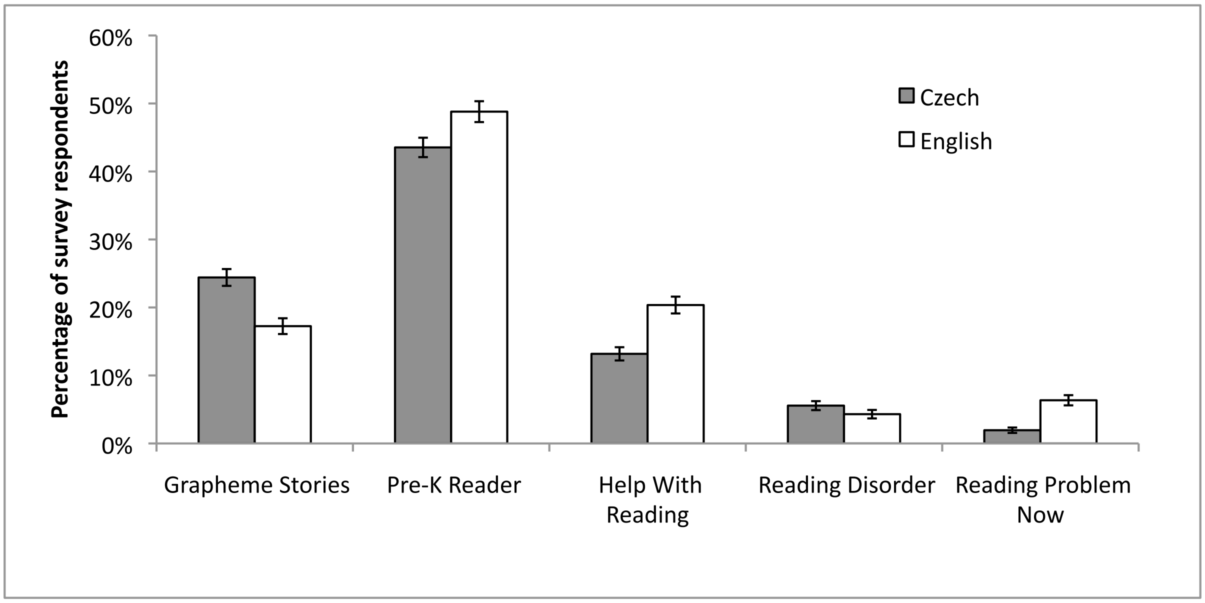

Figure 2.6 Endorsement of questions concerning childhood reading and tendencies to tell stories involving graphemes, in native Czech and English speakers. Error bars indicate +/- 95% C.I.s.

Positive responses to several of these questions were strongly associated with reports of synaesthesia, but none with confirmed synaesthesia. Participants who reported telling stories with graphemes as characters were more likely to report all forms of synaesthesia in both the CU and SFU samples. Reports of learning to read before kindergarten or having reading difficulties, on the other hand, were also associated with reports of most forms of synaesthesia, but only in the SFU sample.

These qualitative descriptions were supported by a series of Fisher's exact tests. For each variety of synaesthesia and each question about grapheme stories or reading, we constructed a 2-by-2 contingency table, where one dimension separated those who endorsed the variety of synaesthesia (or were confirmed as having it on the Battery) from those who did not, and one separated those who endorsed the question from those who did not. Due to very low cell values, the three questions concerning reading difficulty were combined into a single variable: anyone who reported having at least one of these difficulties was designated as having indicated a reading problem.

This left one question asking about grapheme stories, one about learning to read before kindergarten, and one about reading problems in general. Endorsements of 8 varieties of synaesthesia included on the survey were compared to each of these 3 questions, and so *p*-values were Bonferroni-corrected by multiplying them by 24. All forms of synaesthesia were positively associated with reporting having told stories involving graphemes as a child (all *p*s < .001 in both samples), save for the open-ended question about "other" types of synaesthesia which was not significant even prior to Bonferroni correction in the CU sample (*p* > .1), but was marginally significant after correction in the SFU sample (*p* = .06). Reports of early reading were associated with reports of all forms of synaesthesia in the SFU sample (all *p*s < .05) save for number-forms and "other" varieties (both *p*s > .9), whereas in the CU sample early reading was significantly associated with sound-colour synaesthesia (*p* < 001), and marginally associated with reports of letter-colour (*p* = .07), number-colour (*p* = .08) and word-taste (*p* = .08) synaesthesia. Reports of reading problems were associated with all forms of synaesthesia in the SFU sample (all *p*s < .05), but none in the CU sample (all *p*s > .2). None of these questions were associated with confirmed synaesthesias in either sample (all *p*s > .1). (It should be kept in mind that since there are far fewer confirmed than reported synaesthetes, power is clearly an issue here.)

The survey also included two questions concerning gender (CU: 69% female, SFU: 63%) and handedness (CU: 89.4% right-handed, SFU: 87.3%). The associations between gender and synaesthesia are described in the following section. Handedness was not associated with any form of synaesthesia, reported or confirmed (all *p*s > .1 before correction).

2.3.6 A female bias for synaesthesia due to differences in compliance

Women were both more likely to report synaesthesia and to be confirmed as having synaesthesia, a trend which was true for virtually all varieties of synaesthesia in the study, but was much stronger among CU than SFU participants. However, the bias in confirmed rates was almost entirely due to women being more likely to comply with our request to take part in the second part of the study, suggesting that response and compliance biases were the true source of the sex differences we found.

Table 2.1 shows the female:male bias for each of the forms of synaesthesia reported on the paper survey, while Table 2.2 shows the bias for forms of synaesthesia confirmed by the battery. In general, there was a far stronger female bias among CU participants, with only two reported rates and one confirmed rate showing a significant bias among SFU participants, although in almost every case there was a non-significant trend towards one.

Table 2.1 Female:male relative rates of reported synaesthesia

	Letter-Colour	Number-Colour	Weekday/Month-Colour	Sound-Colour	Word-Taste	Number Forms	Grapheme-Personality	Other
Czech	1.50***	1.55***	1.49***	1.13 .	1.24***	1.08	1.18*	0.95
English	0.97	1.05	1.22*	1.09	1.08	1.07	1.21**	1.08

P-values are computed using Fisher's exact method. All *p* values are Bonferroni corrected and > .1, except: . $p < .1$, * p < .05, ** $p < .01$, *** $p < .001$

Table 2.2 Female:male relative rates of confirmed synaesthesia

	Letter-Colour	Number-Colour	Weekday-Colour	Month-Colour	Chord-Colour	Instrument-Colour	Scale-Colour
Czech	3.51*	2.32 .	3.13***	2.21**	0.45	3.31**	3.13
English	3.09*	1.19	1.85	1.68	0.20	4.16	0.59

P-values are computed using Fisher's exact method. All *p* values are Bonferroni corrected and > .1, except: . $p < .1$, * p < .05, ** $p < .01$, *** $p < .001$

The degree of female bias increased sharply from reported to confirmed rates for almost every form of synaesthesia in both groups. One possible explanation is that there was a higher rate of false reporting of synaesthesia among males, or at least a higher proportion of men who reported synaesthetic experiences that are not consistent over time and hence not confirmable by the Synesthesia Battery. Another possibility, however, is that men were less likely to comply with Phase 2 of the experiment, meaning that they would never have been tested in the first place.

In order to test this we examined the female:male relative rates at various stages of the experiment. Specifically, we looked at rates of those who reported synaesthesia, of those who registered for the Synesthesia Battery after being invited, of those who completed at least one test on the Battery, and finally of those who were confirmed synaesthetic by the Battery. Since the female bias was present in both samples and across virtually all forms of synaesthesia, all 11,664 participants were used in order to maximize power (results were qualitatively identical when run with either sample alone), and a single measure was used to indicate if a participant reported or was confirmed as having any one of the seven varieties of synaesthesia tested on the Battery, as in the analysis of second-language acquisition shown in Figure 2.4. Confidence intervals for relative rates were calculated using an online relative risk calculator (MedCalc Software, 2013).

Figure 2.7 shows the results. There was a small female:male bias (1.28) for reporting synaesthesia, then a larger bias (1.67) for registering for the Battery after reporting synaesthesia. No significant bias existed for either completing a Battery test after registering for the Battery (1.09) or being confirmed as synaesthetic after completing a test (1.01). This means that the increase in female:male bias between reported and confirmed synaesthesia was entirely due to attrition: men were less likely to comply with our request to register for the Battery. Of those who did register for the Battery, equal rates of men and women were confirmed synaesthetic. Thus we have no evidence for a female bias in synaesthesia beyond a mild difference in initially reported rates.

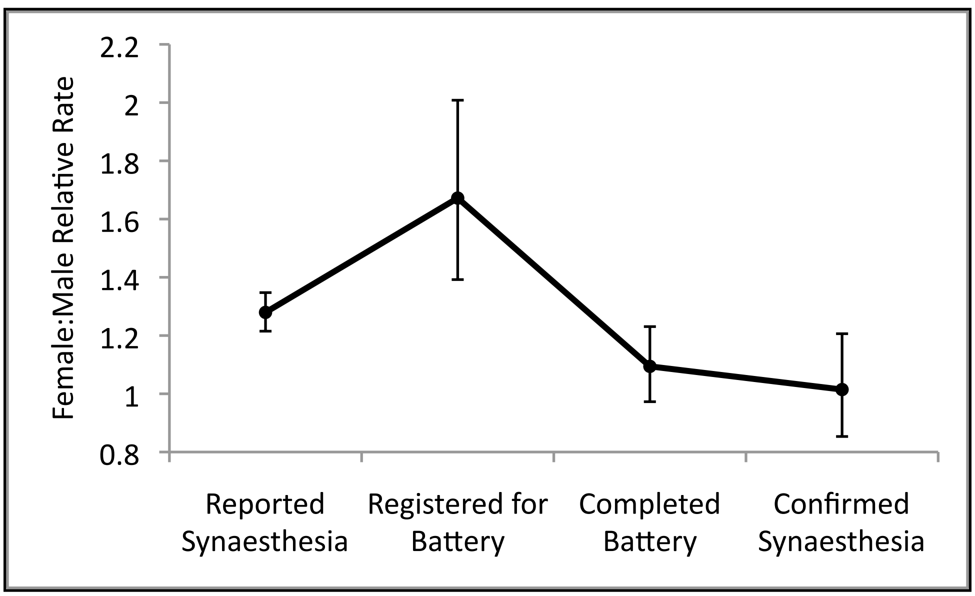

Figure 2.7 Female:male relative rates across various stages of the experiment. There is female bias for reporting synaesthesia and registering for the Synesthesia Battery, but almost identical female:male rates of completing the Battery and being confirmed as synaesthetic. Error bars are +/- 95% confidence intervals

2.3.7 All forms of synaesthesia cluster together

Virtually every form of synaesthesia, whether reported in Phase 1 of the study or confirmed in Phase 2, was associated with every other form of synaesthesia, in both the CU and SFU populations. In other words, individuals with only one variety of synaesthesia, either reported or confirmed, were extremely rare in these data.

This qualitative description is supported by a series of Fisher's exact tests. First, each of the eight forms of synaesthesia asked about on the survey (including "other" as a single type) was compared to each of the 7 other forms, by constructing a 2-by-2 contingency table, where each dimension separated those who endorsed one form of synaesthesia from those who did not. As there were 56 total tests, *p*-values were Bonferroni-corrected by multiplying them by 56. All *p*-values were still significant below the .001 level, save for the comparison between endorsements of grapheme-personality and "other" types of synaesthesia in the SFU sample, where $p = .004$. All these differences were due to a positive association, such that the relative rate of reporting one type of synaesthesia was 1.5-4.9 times higher among those who had reported another type than it was among those who did not report any other type. There were no apparent differences between the CU and SFU samples in these results. (See Appendix 2 for a complete table of relative rates.)

Results were similar for the seven varieties of synaesthesia confirmed on the Battery. Here there were 42 total tests, and so *p*-values were Bonferroni-corrected by multiplying them by 42. Letter-, number-, weekday-, month, and instrument-colour synaesthesias were all positively associated with each other in both the CU and SFU samples (all *p*s < .001), chord-colour synaesthesia was positively associated with only month-colour synaesthesia in both samples (both *p*s < .05), while scale-colour synaesthesia was positively associated with letter-, number-, and instrument-colour synaesthesia in the CU sample (all *p*s < .01) and with letter- and month-colour synaesthesia in the SFU sample (both *p*s < .05). Effect sizes were frankly enormous, with relative rates for the significant comparisons ranging from 23.9-185.9. Given that scale-colour and chord-colour synaesthesia were confirmed in such tiny percentages of both samples (see Figure 2.3), it seems likely that a failure to find positive associations with some other varieties of synaesthesia was simply a power issue. Further evidence in support of this claim comes from the fact that of 8 non-significant comparisons in the CU sample and 9 in the SFU sample, 4 and 6 (respectively) were significant prior to the Bonferroni corrections, and there was no overlap between the 4 remaining non-significant comparisons in the CU sample and the 3 in the SFU sample. (See Appendix 3 for a complete table of relative rates.)

2.4 Discussion

2.4.1 Overview of results

The prevalence of synaesthesia differs markedly between our samples of native Czech and English speakers, with Czechs reporting and being confirmed as having higher rates of almost all forms of synaesthesia—an overall prevalence of 3.7% synaesthetes as opposed to 2.3% among our native English speakers. However this difference does not stem from differences in orthographic transparency, or any other difference between the Czech and English languages *per se*. Rather, it is tied to second language acquisition: those who learn a second language as a native speaker (before age 2) are three times *less* likely to be synaesthetes as those who learn a second language later in life, and there are far more native bilinguals found in the SFU sample. Reading ability may be another influence on the development of synaesthesia, but this result is far stronger among our English than Czech speakers, and is only measurable in rates of reported, not confirmed, synaesthesia. We also find a female bias among our reported and confirmed synaesthetes, however further analysis demonstrates that the confirmed bias is almost entirely due to differential rates of attrition between men and women: men are less likely than women to participate in the second phase of our study. Finally, all forms of synaesthesia are highly associated, such that having one type means one is much more likely to have any other type.

2.4.2 Rates of synaesthesia are consistent with previous studies, likely under-

estimates of true population rates

The confirmed rates of synaesthesia established here are within the general range of those from the two previous studies that use a large sample not dependent on self-referral and use a formal test of synaesthetic consistency (Rothen & Meier, 2010b; Simner et al., 2006). More specifically, four types of synaesthesia that we determine rates for (letter-, number-, weekday-, and month-colour) were also tested for in the University study of Simner *et al*. (2006), and both the CU and SFU rates for these four types are all within the 95% confidence intervals of this study's results. Thus it is clear that Simner *et al*. (2006) were correct to argue that appropriate methodologies are the key to consistent results across different studies. By avoiding the liberal bias of only using the self-report criterion as a test of synaesthesia, and the conservative bias of depending on self-referral to obtain subjects, a clearer picture of the epidemiological profile of synaesthesia can be established. (The fact that both SFU and CU rates are within these confidence intervals, however, also establishes that previous studies lacked the power to find the relatively large group effects we find in the present study.)

While our study's large sample size allows for a high degree of confidence in the results, the proportion of confirmed synaesthesia should be interpreted as lower bounds on its true prevalence, rather than estimates of the actual rate. There are two reasons for this. First, attrition or non-compliance at various stages of the experiment meant that were undoubtedly some genuine synaesthetes who were never tested. A large number of survey respondents did not provide an email address (CU: 31.4%, SFU: 7.1%), meaning that they could not be contacted to participate in the second phase of the study. Of those who were contacted, only approximately one quarter registered for the Synesthesia Battery (CU: 25.6%, SFU: 21.2%), and of those who registered for the Battery, many did not complete a single consistency test (CU: 19.4%, SFU: 27.1%), even though we used a very liberal criterion for "completing" a test: participants simply had to assign any colors at all to at least one inducer on each of the three occasions that this inducer was presented to them during the test.

True synaesthetes would likely be far more interested in the study, and thus more willing to comply with its various stages. Some evidence in favour of this comes from the fact that the rate of endorsing some form of grapheme-color synaesthesia was far lower among participants who did not provide an email address (CU: 13.1%, SFU: 9.2%) than among those who did (CU: 24.8%, SFU: 14.5%), and the same pattern holds for all other forms of synaesthesia on the survey. We have followed up informally with a number of the participants who responded positively to questions about synaesthetic experiences on the survey but did not register for or complete the Battery, all of whom indicated that they misunderstood the initial question or that they were referring to one-off experiences rather than consistent experiences across their lifetime. This is consistent with numerous anecdotal reports from other synaesthesia researchers, as well as the finding of Simner *et al*. (2006) that the majority of individuals who respond positively to questions about synaesthetic experiences are unable to successfully complete consistency tests. Furthermore, we have already noted that our results are within the ballpark of previous studies, in particular our rates for grapheme-colour synaesthesia are close to those established in the Museum study of Simner *et al*. (2006) and in the control study of Rothen & Meier (2010b), both of which were designed such that attrition was not an issue (all participants completed a consistency test). Nevertheless, there are almost certainly some genuine synaesthetes who, for whatever reason, did not provide an email address, or did not register for or complete the Battery.

The second reason why our rates should be interpreted as lower bounds stems from the fact that we found a surprisingly large number of confirmed synaesthetes who reported on the survey that they did not have the form of synaesthesia they were confirmed as having. This raises important questions about the appropriate definition and operationalization of synaesthesia, as clearly the self-report criterion that is used by almost all studies of synaesthesia is not only prone to false positives (in that most individuals who meet the criterion do not meet the consistency criterion) but also to false negatives (in that some people who do not meet the criterion *do* meet the consistency criterion). A more pressing concern for interpreting these results, however, is that it is unclear how many of those who report no synaesthetic associations at all would be able to make consistent associations. We see no reason to think that this group is insignificant.

Clearly the best way of solving these issues would be to run a similarly-sized study where *all* participants were given consistency tests, regardless of their self-reported experiences. For the time being, we will simply have to accept that the true rates of synaesthesia are likely somewhat higher than those established in this study and in the University study of Simner et al (2006), which also only gave consistency tests to those who reported synaesthesia.

2.4.3 Learning and synaesthesia

One of the hypotheses driving this study was that the orthographic opaqueness of English would result in more English than Czech grapheme-colour synaesthetes, but not affect other varieties of synaesthesia. Clearly, this hypothesis is false, both in terms of the direction of the effects and their specificity. Czechs, not English speakers, are more likely to be synaesthetic, and this is true across synaesthesia in general, not restricted to the grapheme-colour variety. And in perhaps the most intriguing finding from the present study, we find that this group difference is associated with second-language learning: those who learn a second language later in life are more likely to develop synaesthesia, of virtually any type. This single factor appears to entirely account for the differences between our two language groups.

This result was, to say the least, a surprise. Bilingualism is associated with stronger executive functioning and a higher degree of metalinguistic awareness (Bialystok & Barac, 2012), which might either lead to greater or lesser rates of grapheme-colour synaesthesia (supporting the simplicity and complexity hypotheses, respectively). Indeed, our initial reason for asking about the age of acquisition of second languages was simply to rule out those who learned a second language as an adult. However there is no significant difference between the rates of synaesthesia between bilinguals and monolinguals, rather the difference is between *native* and *non-native* bilinguals.

Unfortunately, any attempts to explain this finding will suffer from a serious lack of hard data. Most research on bilingualism tends to be performed on non-native bilinguals, and most research on native bilinguals compares their performance to monolinguals, not to non-native bilinguals (cf. Werker & Byers-Heinlein, 2008). The research that does exist, in a nutshell, shows that non-native bilinguals are faced with a harder task than native bilinguals, and that they have a number of corresponding neurophysiological and behavioral differences. In general, the brain areas activated by the use of second languages do not differ with the age of acquisition, but children who acquire their second languages later show more variable and greater activation in these areas (Bloch et al., 2009), and are less able to suppress the involuntary switching of attention (Ortiz-Mantilla, Choudhury, Alvarez, & Benasich, 2010). Furthermore, if participants are young enough, then studies comparing monolinguals and native bilinguals are effectively also comparing non-native bilinguals, since some of the monolinguals will go on to learn a second language later in life. There are several reliable differences of note here. In a nutshell, young native bilinguals tend to have a smaller vocabulary in each of their two languages (e.g. Bialystock & Herman, 1999; Cobo-Lewis, Pearson, Eilers, & Umbel, 2002), a greater degree of phonological awareness (Bialystock, Luk, & Kwan, 2005), and stronger executive functions (Kovács & Mehler, 2009).

Non-native bilinguals are, by definition, faced with a difficult learning task that monolinguals do not face: namely learning a new language, and more specifically learning to read and write in this new language. Native bilinguals will generally also become literate in both their languages, but they have the distinct advantage of already being fluent speakers of each language, and also a higher degree of phonological awareness and executive function that will likely assist in becoming literate. Thus learning to read and write in both languages may be significantly easier for native than non-native bilinguals, and since monolinguals do not have a second language to learn, this allows them to be differentiated from both bilingual groups. This explanation, then, is consistent with the complexity hypothesis and not with the simplicity hypothesis. However more research is clearly needed in this case.

2.4.4 Development or retention of synaesthesia

An important question is whether the bilingualism effects we observe are due to non-native bilinguals being more likely to develop synaesthesia, or less likely to lose it. This chapter has generally advanced the idea that differences in adult rates of synaesthesia are due to differences in developing synaesthesia, but there is no *a priori* reason why these effects could not arise due to differences in retaining an already-existing synaesthesia. The idea that the utility of synaesthesia might cause either its development or simply its retention is an old one (Calkins, 1893), and the latter possibility is consistent with the notion that synaesthesia is a normal part of development (Maurer, 1993). The present data, of course, do not allow these two possibilities to be decided.

2.4.5 The generalizability of synaesthetic tendencies

We initially hypothesized that group differences would only be found for grapheme-colour synaesthesia, but instead found a general difference between groups for almost all varieties of synaesthesia, and a strong tendency for all types of synaesthesia that we test for to cluster with each other. This is consistent with Novich *et al.*'s (2011) reports of clusters of synaesthetic types, since almost all the varieties of synaesthesia confirmed on the Battery are what they refer to as coloured-sequence synaesthesia. Novich *et al.* (2011) provide a neurological explanation of this clustering, suggesting that coloured-sequence synaesthesias are based in unusual connectivity found within the neural networks responsible for coding sequence information, located in the middle temporal gyrus and temporoparietal junction in the right hemisphere, and the inferior frontal gyrus in the left hemisphere (Pariyadath, Plitt, Churchill, & Eagleman, 2012). We would like to offer another, learning-based, hypothesis for this clustering: perhaps once synaesthesia has been employed in solving a given task, it is likely to be used for other, conceptually similar tasks. Of course this is not meant as an alternative to the neurological account of Novich *et al.* (2011), and both factors might play an important role.

2.4.6 Sex bias in synaesthesia

Our results firmly establish that women are more likely than men to report synaesthetic experiences, and to voluntarily follow experimental procedures testing for synaesthesia (at least in the North American and European contexts in which these studies have taken place). It is technically possible that these differential rates of attrition and non-compliance between men and women are due to more men realizing that their false reports of synaesthesia will be uncovered by rigorous consistency testing, and thus feeling too uncomfortable to proceed. However this seems overly complex, especially in the face of the alternative explanation that women in our societies are simply more likely to follow requests and instructions, particularly when these involve self-disclosure of things that may be embarrassing or at least uncomfortable (cf. Dindia & Allen, 1992; Simner et al., 2006; Ward & Simner, 2005).

Like the two previous studies that have shown a smaller sex bias in synaesthesia than was previously thought to exist (Simner et al., 2006; Ward & Simner, 2005), our findings show a small female bias that cannot be explained on the basis of attrition (an overall relative ratio associated with being female of 1.28 for reporting synaesthesia). It is possible that this reflects a genuine sex difference in prevalence, but once more it seems far more parsimonious to explain this bias on the basis of a self-disclosure bias: male synaesthetes may be less likely than women to report synaesthesia *even when asked directly about it*. The only appropriate way of falsifying this hypothesis, once again, is a large-scale study that requires all participants to take a consistency test regardless of self-report. Until such time, we can be confident based on these results that if there is a genuine sex bias in synaesthesia, it is a minor one.

If men are less likely to report synaesthesia even when asked directly about it, this casts doubt upon reports that synaesthesia is predominantly passed down via the maternal line (Barnett et al., 2008a; Baron-Cohen, Burt, Smith-Laittan, Harrison, & Bolton, 1996; Ward & Simner, 2005). All these studies screened participants using a short questionnaire, and did no further testing on those who responded negatively. Thus it could very well be that there are far more synaesthetic fathers of synaesthetes than previously thought.

In general, then, these results strongly imply that there is no true sex bias in synaesthesia, which further weakens one of the original motivations for assuming a simple genetic cause of synaesthesia.

2.4.7 Problems with the self-report and consistency criteria?

Two aspects of our results raise questions about the suitability of the self-report criterion for synaesthesia. First, as suggested above, it is very likely the case that men are less likely to report synaesthesia even when asked directly about it. Second, a large fraction of our confirmed synaesthetes denied having the type of synaesthesia on the initial survey which they were eventually confirmed to have by the consistency criterion. Both these results indicate that self-reported synaesthetic experience is not necessary to have highly consistent inducer-concurrent relationships. Conscious experience has typically been taken as a fundamental aspect of synaesthesia, but we suggest that this is not necessarily true.

Furthermore, there may be reason to think that the consistency criterion is not an appropriate test for all forms of synaesthesia. Roughly the same proportion of people report colour experiences induced by letters and by music, but the consistency test is passed by far more of those who report letter colours than music colours. It is possible that individuals interpret questions about music-induced colours in a more metaphorical manner than questions about letter-induced colours, and thus that questions about music-colour synaesthetes will be answered far more liberally. However it seems equally plausible that music tends to induce colour experiences that are highly context-dependent, and thus not amenable to simple consistency tests. This may be the real explanation for the very low rates of confirmed music-colour synaesthesia in this and other studies (e.g. Novich, Cheng, & Eagleman, 2011).

2.4.8 Future directions

The epidemiology of synaesthesia is still somewhat mysterious. There are strong indications that the standard criteria used to confirm synaesthesia have important flaws. It also appears that in order to firmly establish both the precise rates of synaesthesia in the population and the nonexistence of a female bias among synaesthetes, a large-scale study that does not screen participants based on self-report will be necessary. And new research will also be needed to confirm or refute the hypothesis that the greater rates of synaesthesia among non-native bilinguals are due to the increased demands they face while becoming literate in their second language, and to determine if these greater rates are due to the development or retention of synaesthesia. However our results have shed light on some mysteries. Most importantly, they provide a strong reason to think that the nature of the learning challenges faced in childhood have a large impact on the development of synaesthesia.

3 Second-order mappings in grapheme-colour synaesthesia[2]

3.1 Introduction

Despite an explosion of research on grapheme–colour synaesthesia over the past two decades, little is known about how these associations are made. Why does Jane see the letter M as a deep purple, while John associates the same letter with forest green? Here we verify that there are several different sources of synaesthetic associations, and we investigate both how they interact with each other and what aspects of synaesthetic colour they influence.

To date, synaesthesia research has documented a number of regularities in the grapheme–colour pairs of individuals. For example, English speakers often associate the letter B with blue or brown, G with green, and so on for the first letters of other common colour names (Barnett et al., 2008a; Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2005). Similarly, some synaesthetes have adopted the colours of letter-shaped fridge magnets used in their childhoods (Witthoft & Winawer, 2006; Witthoft & Winawer, 2013). These are regularities in *first-order relations*—that is, between nonrelational

2. This chapter is a slightly adapted version of a previously published paper (Watson, Akins, & Enns, 2012). The data for this paper was kindly provided by Michael Dixon and Jonathan Carriere of the University of Waterloo. The initial research question was collaboratively arrived at by the three co-authors, all analyses were my own, and I was the principal author of the paper.

properties of a letter (such as its shape or name) and dimensions of synaesthetic colour such as hue and lightness (Day, 2005).

A parallel line of research has begun to investigate grapheme–colour pairings by looking for *second-order relations*, or “relations between relations.” For example, letters with similar shapes, such as E and F, tend to be associated with synaesthetic colours that are similar in hue (Brang, Rouw, Ramachandran, & Coulson, 2011; Eagleman, 2010; Jürgens & Nikolic, 2012; Brang, Rouw, Ramachandran, & Coulson, 2010).

Here there is a correlation between two relations: A relation of similarity in the domain of letter shape is correlated with a relation of similarity in the domain of synaesthetic colour. Importantly, second-order relations can exist independently of first-order pairings. That is, two synaesthetes may each assign different colours to E, but so long as each individual’s colour for F is similar to that individual’s colour for E, this constitutes a second-order relation between letter shape and synaesthetic colour. Thus, second-order letter–colour associations may not be apparent when looking at first-order relations.

A variety of second-order influences on synaesthetic colour have been demonstrated. Marks (1975) noted that music-colour synaesthetes often associate higher pitches with brighter colours. In grapheme-colour synaesthesia, numerals and letters that appear more frequently in print tend to be associated with brighter (Beeli, Esslen, & Jäncke, 2007; Cohen Kadosh, Henik, & Walsh, 2007; Simner & Ward, 2008; Smilek, Carriere, Dixon, & Merikle, 2007) and more saturated (Beeli, Esslen, & Jäncke, 2007) colours. More frequent letters also tend to be associated with colours whose names are more common in spoken language (Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2005). Each of these results has been reported as a first-order relation (correlations between *absolute values* on two dimensions), but they all imply second-order relations (correlations between *differences in values* on two dimensions). For example, the fact that more-frequent letters have brighter colours implies that letters that differ greatly in terms of frequency will *also* differ in terms of their brightness.

Two recent results have come directly from second-order analyses. First, as noted above, letters with similar shapes appear to be associated with similar synaesthetic colours (Brang, Rouw, Ramachandran, & Coulson, 2011; Eagleman, 2010; Jürgens & Nikolic, 2012; Brang, Rouw, Ramachandran, & Coulson, 2010). Second, letters early in the alphabet tend to have colours that are quite distinct from each other, whereas later letters tend to have colours that are more similar to those of nearby letters (Eagleman, 2010). On Eagleman's view, this pattern stems from the order in which children learn their letters. The first letter learned is associated with an idiosyncratic colour; the next letter is associated with a colour that is easily distinguishable from the first; and each subsequently learned letter is associated with a colour as distinct as possible from those already assigned. With each letter learned, however, the range of distinct colour choices is diminished, and inevitably, letters learned later are associated with colours similar to some of those associated with earlier letters. Note that this interpretation implies a relation between letter ordinality and synaesthetic colour that is similar to Weber's fraction. In brief, a pair of letters that appear early in the alphabet (e.g., A and D) will be assigned colours that are more distinctive than will a pair of letters later in the alphabet (e.g., S and V), even though they are equal numbers of steps apart in absolute units (three, in this example). Such a finding requires a second-order perspective: When one looks at absolute hue assignments, no relation with ordinality is found (Cohen Kadosh, Henik, & Walsh, 2007; Smilek, Carriere, Dixon, & Merikle, 2007).

In line with these findings, we prefer to analyze second-order relations among grapheme–colour pairs. Our primary motivation is that strong second-order mappings (with weaker first-order mappings) have often been observed in human perception more generally (e.g., we remember melodies, not absolute pitch, in music; facial configurations, not specific facial features, in vision; and words, not phonemes, in language). A secondary motivation is that second-order analyses allow for the easy investigation of the property of hue. Because luminance and saturation are one-dimensional properties of colours, they can be used in correlations or other linear analyses. Hue requires at least two dimensions, however, in order to be specified (e.g., blue-yellow or red-green), which makes it impossible to compute a simple correlation between hue and any other measure. *Differences* between hues, on the other hand, are one-dimensional, and thus amenable to linear analysis.

In the present study, we compared the colours assigned to letters by a large group of synaesthetes (N=54) with a wide variety of letter similarity measures taken from non-synaesthetic individuals. We sought to determine how different aspects of letter similarity (e.g., shape, order, and frequency) are related to synaesthetic colours, and how these effects relate to each other. We were especially interested in how the various aspects of letter similarity might be related to two dimensions of colour—namely, luminance and hue (Beeli, Esslen, & Jäncke, 2007). As noted above, differences in letter frequency have been shown to correspond to differences in luminance, while differences in hue have generally been overlooked, possibly because researchers have been looking for first-order relations. How letter shape and ordinality map separately onto luminance and hue remains an open question.

3.2 Data preparation

The RGB colour values of each letter were provided by 54 confirmed grapheme–colour synaesthetes (Smilek, Carriere, Dixon, & Merikle, 2007). These values were recoded into CieLab colour space, which more accurately describes human colour discriminations and allows for the separation of colour into luminance and hue components. There are 325 possible letter pairs (not including doubles of the same letter), and for each of these pairs we computed separate values for colour distance (Euclidean distance in CieLab space), luminance distance (distance along the CieLab L-axis), and hue distance (distance in the CieLab *ab* plane). These values were averages of the distances across all 54 synaesthetes.

Table 3.1 Letter similarity measures used in the English studies

Similarity Measure	Description
Shape difference	Euclidean distance in an 11-dimensional space defined using the basic letter shape features from Gibson (Gibson, 1969)
Frequency difference	Difference of two letters' frequencies divided by the sum of their frequencies (Lewand, 2000)
Ordinality difference	Difference of two letters' positions in the alphabet divided by the sum of their positions
Ratings	
A (similarity)	Similarity ratings of uppercase letters (Boles & Clifford, 1989)
B (similarity)	Similarity ratings of lowercase letters (Boles & Clifford, 1989)
C (difference)	Difference ratings of uppercase letters (Podgorny & Garner, 1979)
Discrimination RT	Reaction time on a same–different discrimination task for uppercase letter pairs (Podgorny & Garner, 1979)
Confusion	
A	Chance of confusing two briefly presented uppercase letters on a letter-naming task (Gilmore, Hersh, Caramazza, & Griffin, 1979)
B	Chance of confusing two uppercase letters (in Keepsake font) presented at low intensity on a letter-naming task (Gupta, Geyer, & Maalouf, 1983)
C	Chance of confusing two uppercase letters (dot-matrix font) presented at low intensity on a letter-naming task (Gupta, Geyer, & Maalouf, 1983)
Letter Name Similarity	Number of shared phonemes in two letter names (e.g., "bee" and "dee" have 1 shared phoneme, /i/)

A total of 11 measures of letter similarity were derived for comparisons with the synaesthetic colour data (see Table 3.1). *Shape difference* is the Euclidean distance in a letter-shape similarity space generated from 11 basic letter-shape features (Gibson, 1969), such as the presence or absence of a diagonal line (see Table 3.2). *Frequency difference* and *ordinality difference* are the differences between the frequencies (Lewand, 2000) and positions in the alphabet of two letters, divided by their sum. *Letter name similarity* consists of the number of shared phonemes in the English names of two letters; for instance, the names of the letters B and D share one phoneme, /i/, and hence would have a letter name similarity of 1 (Ward & Simner, 2003). These are examples of the familiar Weber fraction that describes perceived difference in numerous psychophysical domains. The remaining measures were previously published behavioral data on letter similarity, and thus may have been influenced by letter shape, frequency, order of acquisition, and (potentially) many other factors. These measures include *discrimination RTs*, from a same–different task in which the subjects were briefly presented with letter pairs (Podgorny & Garner, 1979); *comparison ratings* of letter similarity or difference (Boles & Clifford, 1989; Podgorny & Garner, 1979); and *confusion*, from letter-naming tasks using degraded stimuli (Gilmore, Hersh, Caramazza, & Griffin, 1979; Gupta, Geyer, & Maalouf, 1983).

Except where noted, all subsequent analyses were performed after binning the 325 letter pairs into 65 bins that each included five letter pairs. Bins were determined by the mean colour distance of each letter pair across all 54 synaesthetes, such that the first bin contained the five pairs whose two letters were, on average, most similar in colour, and the last bin contained the five pairs whose two letters were, on average, most dissimilar in colour.

Table 3.2 Letter shape dimensions adapted from Gibson (1969). Some dimensions' names have been changed for clarity, and two dimensions—right and left diagonal—have been combined into a single diagonal dimension.

Shape Dimension	Letters
Straight Horizontal	*A E F G H L T Z*
Straight Vertical	*B D E F H I K L M N P R T Y*
Closed Curve	*B D O P Q R*
Upward-Opening Curve	*J U*
Horizontal-Opening Curve	*C G J S*
Intersection	*A B E F H I K P Q R T X*
Repeated Element	*B E M S W*
Symmetry	*A B C D E H I K M O T U V W X Y*
Vertical Discontinuity	*A F H I K M N P R T Y*
Horizontal Discontinuity	*E F L T Z*
Diagonal	*A K M N Q R V W X Y Z*

3.3 Results

3.3.1 Letter similarity measures predict different aspects of colour similarity

We computed the simple correlations of all of the letter similarity measures with colour, luminance, and hue distance (see Table 3.3). Since multiple correlations were run, we corrected the p values, multiplying each by 7, as seven distinct types of measures were being compared with each of the colour distance measures. We also used Spearman's rho for the correlations involving the three ratings, since they are ordinal measures.

Colour distance and hue distance were both correlated with shape difference, ordinality difference, and Letter Confusion B. Luminance distance was correlated with frequency difference, with Rating C and marginally correlated with Letter Confusion C. Thus, there appears to be a split between those aspects of letter similarity that predict synaesthetic luminance and hue.

Table 3.3 Correlations between letter and colour similarity

Similarity Measure	Colour Distance	Luminance Distance	Hue Distance
Shape difference	0.48***	0.07	0.49***
Frequency difference	0.06	0.34*	0.01
Ordinality difference	0.37*	0.02	0.39**
Ratings			
A (similarity)	-0.20	-0.28	-0.19
B (similarity)	-0.22	-0.29	-0.20
C (difference)	0.23	0.35*	0.21
Discrimination RT	-0.21	-0.26	-0.19
Confusion			
A	0.04	-0.23	0.07
B	-0.34*	-0.26	-0.32 .
C	-0.27	-0.31 .	-0.24
Letter Name Similarity	-0.02	-0.00	-0.02

Correlations with the ratings use Spearman's rho. All p values are Bonferroni corrected and > .1, except: . $p < .1$, * $p < .05$, ** $p < .01$, *** $p < .001$

3.3.2 *Letter shape and ordinality predict hue; letter frequency predicts luminance*

All of the correlations described above can be accounted for in terms of only three mappings, shown in Figure 3.1. A first mapping involves letter shape and synaesthetic hue, a second involves letter ordinality and hue, and a third involves letter frequency and luminance. A regression model using only shape difference and ordinality difference to predict hue distance (R^2 0 .31, $p < .001$) did not explain less variance than one using all 11 letter similarity measures to predict hue distance ($p > .05$). However, removing either shape difference or ordinality difference from the reduced model resulted in significantly less explained variance ($p < .05$ in both cases). As Confusion B was also significantly correlated with hue distance (see Figure 3.3), we tried adding it to this reduced model, but it did not explain any variance independently of shape and ordinality difference ($p > .9$). Similarly, a regression model using frequency difference as the sole predictor of luminance distance (R^2 0 .12, $p < .01$) did not differ from a model using all 11 similarity measures as predictors ($p > .1$). We also tried a two-predictor model that included Rating C, as this was correlated with luminance distance, but it did not explain any variance independent of frequency difference ($p > .1$).

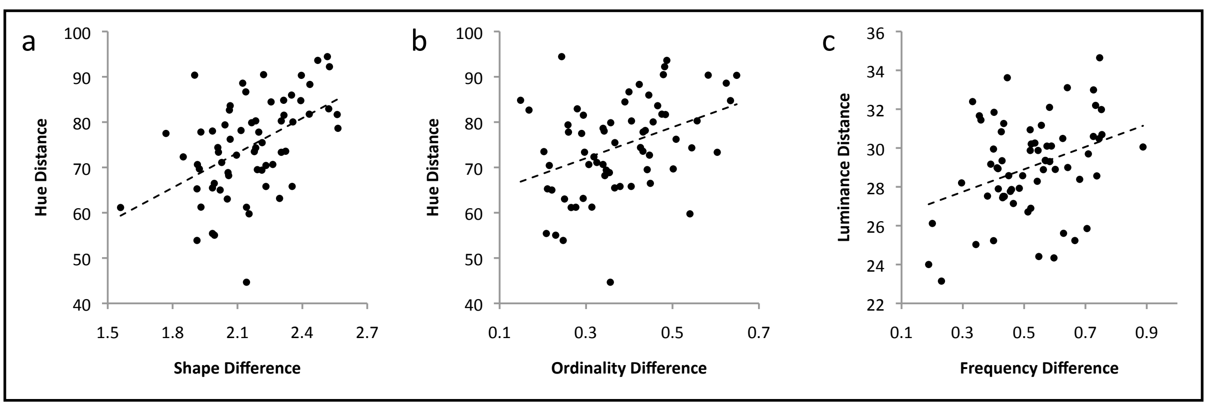

Figure 3.1 Scatterplots of three second-order mappings between letter similarity and synesthetic color. The x-axes denote differences between letter pairs in (a) letter shape, (b) letter ordinality, and (c) letter frequency. The y-axes for panels (a) and (b) denote distance in terms of synesthetic hue; in panel (c), the y-axis denotes distance in terms of synesthetic luminance. The 65 data points in each plot were obtained by binning 325 letter pairs into five-pair bins and then averaging over 54 synesthetes (from Smilek et al., 2007)

3.3.3 Analyses of individual differences show that the mappings are independent

We computed the correlations for each of the mappings in Figure 3.1 at the level of individual synaesthetes. This revealed that these correlations were positive for a majority of the synaesthetes (81%, 67%, and 54% for panels (a), (b), and (c), respectively, in Fig. 1]) Overall, 28% of the synaesthetes had positive correlations for all three mappings, 46% had positive correlations for two of the mappings, 19% had a positive correlation for only one mapping, and the remaining 7% had no positive correlations for any of the three mappings. Critically, there were no hints of correlations between any of these mappings, as tested by coding the presence or absence of each mapping as 0 or 1 for each synaesthete, or by correlating the rank order of synaesthetes on each mapping, as determined by the magnitude of their individual correlations (all ps > .2).

3.3.4 Which aspects of shape matter?

The shape difference measure was further subdivided into 11 dimensions of shape that are important in letter identification (Gibson, 1969). Only two of these dimensions were significantly correlated with hue distance (after a Bonferroni correction)—namely, distance along the closed curve and repeated element dimensions (rs = .41 and .38, ps = .01 and .02, respectively; all other ps > .1). However, a model that used only these two dimensions to predict hue distance explained less variance than did a model using all 11 dimensions ($p < .05$). In the complete 11-predictor model, the only variables that made a significant independent contribution were distance along the closed curve, repeated element, and diagonal dimensions (all ps < .01). Thus, we tried a model using these three variables to predict hue distance ($R^2 = .38$, $p < .001$), and found that it did not predict less variance than the complete 11-predictor model ($p > .1$), but removing any one of these three dimensions from the model resulted in less explained variance (all ps < .05). Further study will be needed to determine why these features are especially important to synaesthetes.

3.4 Discussion

These results confirm that three distinct aspects of letter similarity have a second-order influence on synaesthetic colour assignments. The shape, frequency, and ordinality of individual letters influence the colours assigned to them by synaesthetes, and these three effects are completely independent of each other: For instance, an individual with a strong shape-hue association may or may not have a strong frequency-luminance association. Finally, each of these mappings is confined to a particular dimension of colour space: Letter shape and ordinality are associated with hue, while frequency is associated with luminance.

Brang et al.'s (2011) cascaded cross-tuning model of synaesthesia states that shape–colour associations are the result of the coactivation of contiguous brain areas in the fusiform gyrus that represent letter form and colour. This model does not currently account for the other relations we found, nor for the fact that each relation is confined to a particular dimension of colour. Instead of looking for an explanation at the level of shared neurons, we offer two complementary hypotheses for these findings, both of which revitalize an old hypothesis of Calkins (1893): that synaesthetic associations may arise for strategic reasons. (We remain agnostic as to whether those who employ such strategies are consciously aware of doing so.)

First, associating letter shapes and identities with hue might aid learning to read, but associating them with luminance might compromise reading performance. In vision, a common strategy is to process hue and luminance separately (Gheorghiu & Kingdom, 2006, 2007; Kingdom, Beauce, & Hunter, 2004; Kingdom & Kasrai, 2006; Liebe, Fischer, Logothetis, & Rainer, 2009; Nagai & Uchikawa, 2009; Shimono, Shiori, & Yaguchi, 2009), because each dimension provides different information about the environment (Hansen & Gegenfurtner, 2009). For example, a vital part of vision is to differentiate shadows from material objects. Since shadows are defined by differences in luminance, whereas objects usually differ from their background in both luminance and hue, it follows that hue edges are a more reliable cue to object boundaries than are luminance edges.

A similar moral applies in reading. Graphemes are usually presented as dark, achromatic elements on a lighter background, and thus are usually processed entirely on the basis of luminance contrast. Second-order relations between synaesthetic hue and shape could provide an additional source of information to be exploited for such tasks as letter segmentation, identification, place-holding for visual saccades, search for letters, maintaining letter order in short-term memory, and so forth. However, similar mappings between synaesthetic luminance and shape might interfere with the luminance-sensitive channels responsible for letter shape perception, and so could conflict with natural variations in luminance from the font and from illumination. Thus, synaesthetes may exploit information about letter identity encoded in synaesthetic hue, in addition to the systems that they share with nonsynaesthetes, which use luminance contrast in the various cognitive operations involved in reading.

A complementary hypothesis for mapping hue and luminance to separate aspects of letter identity stems from differences in the ways that humans use hue and luminance to represent information. Take map reading as an example. Defining regions by hue typically allows for faster and more accurate judgments of categorical distinctions than does defining them by luminance, while luminance scales afford advantages for judgments about relative quantity or continuous magnitudes (Breslow, Trafton, McCurry, & Ratwani, 2010). This likely reflects the fact that variations in luminance have an underlying continuity, from dark to light, while hues are perceived categorically. As letter frequency varies along a continuum, then, it maps naturally to luminance. Letter shapes, on the other hand, are perceived categorically (Boles & Clifford, 1989), and thus map naturally to hue. Since letter ordinality also varies continuously, one might think that it should be associated with luminance. Recall, however, that we use letter ordinality as a rough index of the order of learning of individual letters, which are themselves seen as categorical objects (Eagleman, 2010), so the association between ordinality and hue is also consistent with this hypothesis.

Previous research has reported a number of first-order synaesthetic colour associations; for instance, the letters used to begin common colour words are frequently associated with the colours named by these words, and the letters O and I are almost always black, white, or gray (Barnett et al., 2008a; Day, 2005; Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2005). We stress that the finding of second-order relations is complementary to such results. It appears that a wide range of factors, of both the first and second orders, can potentially influence letter–colour mappings, as a good deal of variance in both luminance and hue still remains unexplained. Our analysis of individual differences suggests that these factors often co-exist within individual synaesthetes. That is, any given letter-colour mapping might be influenced by a particular factor, but a different letter (or the same letter for a different synaesthete) is quite likely to be coloured according to a different factor.

In summary, examining relations involving differences between letters and their assigned colours has allowed us to directly compare and contrast multiple influences on synaesthetic associations. The finding that second-order relations are pervasive in synaesthesia is further evidence for the view that synaesthesia builds on normal mechanisms (Barnett et al., 2008a; Simner et al., 2005). Though most of us may not reliably associate letters with colours, those of us who do tend to use principles common to other sensory and cognitive domains.

4 Higher-fidelity synaesthetic colour data increases strength of effects[3]

4.1 Introduction

The previous chapter demonstrated a number of second-order relationships between various properties of letters and their synaesthetic colours. Since analyzing these data, however, I have obtained a similarly-sized database of colours from different synaesthetes. In this chapter I re-run the same analyses as in Chapter 3 on this new data set. This is partly in order to verify that the pattern of results is similar, as Chapter 3 was an exploratory study that found a number of unexpected results, and so one might question if these results generalize. Furthermore, it is probable that the colours in the new data set have a higher fidelity to synaesthetic experience, as they are chosen from all possible colours on a standard computer screen, which vary freely in hue, saturation, and luminance, unlike the colours chosen by participants in Chapter 3, which, in order to speed up colour selection, were all fully saturated. Thus it is possible that there will be stronger effect sizes using the present data set.

4.2 Participants

Potential participants were identified on the basis of their responses to an ongoing large-scale survey about synaesthetic tendencies at Simon Fraser University in British Columbia, Canada, or were self-referred after viewing advertisements on various websites (e.g. Craigslist) or on bulletin boards in the Vancouver area (e.g. at universities, cof-

3. See Footnote 1 on p. 26 for authorship details.

fee shops, libraries, community centres, etc). All potential synaesthetes were invited to take the online synaesthesia Battery (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007), a standard cofirmation test for synaesthesia that verifies that the colour associations made with a variety of inducers (letters, sounds, weekdays, etc) are consistent across repeated presentations. 48 participants obtained a consistency score below 1 on the letter-colour consistency test, which is considered a strong indicator of genuine synaesthesia.

4.3 Data preparation

While completing the Battery, each participant assigned an RGB colour to each letter 3 times. For each letter, we took the final (third) colour assigned to that letter as its canonical synaesthetic colour for that individual, on the assumption that by the third attempt, participants would have fully understood the colour-choosing interface. All analyses use these canonical colours. Participants had the option of choosing "no colour" for particular letters, and if they did this on any trial that letter was ignored in all subsequent analyses for that participant.

As in Chapter 3, all RGB colours were recoded into CieLab colour space, and each participant's 26 letter colours were converted into three different distance measures: *colour distance*, *luminance distance*, and *hue distance*. These values were then averaged across all 48 synaesthetes. The same 11 measures of letter similarity were used for comparisons with the synaesthetic colour data (see Table 3.1). Also as in Chapter 3, all data were binned from 325 letter pairs into 65 bins that each included five letter pairs, ordered according to the mean colour distance across all participants.

4.4 Results

Chapter 3 used four types of analyses. First, simple correlations between the 3 measures of colour distance and the 11 letter similarity measures revealed that luminance and hue were correlated with different aspects of letter similarity. Second, linear models revealed that three effects drove the simple correlations: a relationship between shape and hue, between ordinality and hue, and between frequency and luminance. Third, analyses of individual differences showed that these three effects were independent of each other, such that the likelihood of individuals displaying any of them is unaffected by whether they display any of the others. Finally, by further breaking down shape difference into 11 different dimensions, it was shown that only three of these dimensions (closed curve, repeated element, and diagonal) were related to differences in hue. Each of these analyses is reproduced below.

4.4.1 Shape-hue and ordinality-hue results replicate

Table 4.1 Correlations between letter and colour similarity (new data set)

Similarity Measure	Colour Distance	Luminance Distance	Hue Distance
Shape difference	0.47***	0.21 .	0.45***
Frequency difference	-0.05	0.09	-0.06
Ordinality difference	0.64***	0.02	0.66***
Ratings			
A (similarity)	-0.35**	-0.18	-0.32**
B (similarity)	-0.10	-0.07	-0.10
C (difference)	0.35**	0.23 .	0.31*
Discrimination RT	-0.40***	-0.20	-0.39**
Confusion			
A	-0.14	-0.17	-0.11
B	-0.35**	0.04	-0.36**
C	-0.36**	-0.12	-0.35**
Letter Name Similarity	0.04	-0.20	0.09

Correlations with the ratings use Spearman's rho. All **p** values are Bonferroni corrected and > .1, except: . p < .1, * **p** < .05, ** **p** < .01, *** **p** < .001

Table 4.1 displays the simple correlations between the three dimensions of colour distance and the 11 dimensions of letter similarity in the new data set. A comparison with Table 3.3 on p. 64 reveals a high degree of similarity between the two sets of results. Once again, there were strong correlations between both Colour Distance and Hue Distance with Shape Difference and Ordinality Difference. Furthermore, the overall pattern of positive and negative correlations was preserved quite closely. Indeed, all of the correlations with an absolute magnitude of over 0.1 on Table 4.1 were in the same direction on Table 3.3, and vice versa. In general, the correlations on Table 4.1 are stronger, with twice as many reaching significance (14 as opposed to 7 on Table 3.3). One clear difference between the two data sets, however, is that the correlation between luminance and frequency did not approach significance in the new data set.

A model using shape and ordinality difference to predict hue distance (R^2 = .57, p < .001) did not predict any less variance than a model using all 11 letter similarity measures as predictors (p > .5). However removing either shape or ordinality from the two-predictor model led to less variance explained (p < .001 in both cases). Thus, as in Chapter 3, it seems safe to say that any correlations between the various similarity measures and hue are entirely due to shape and ordinality effects. Furthermore, this model explained almost twice as much variance in hue distance as the same model using data from Chapter 3. The model using all 11 measures to predict luminance difference, on the other hand, barely approached significance (R^2 = .26, p = .09). Thus luminance in the present data set was not well-predicted by the similarity measures, whether we use simple correlations or a linear model.

4.4.2 The shape-hue and ordinality-hue effects are independent

In terms of individual differences, the shape-hue and ordinality-hue correlations were both positive for 81% of participants. However these were not necessarily the *same* participants: 65% of participants showed both effects, while 15% showed only one or the other. As in Chapter 3, there was no hint of a correlation between these mappings, as tested by coding the presence or absence of each mapping as 0 or 1 for each synaesthete, or by correlating the rank order of synaesthetes on each mapping, as determined by the magnitude of their individual correlations (both ps > .2).

4.4.3 The same dimensions of shape predict hue distance

Finally, separating the shape difference measure into distance along 11 separate dimensions led to very similar results as in Chapter 3. Distance along three of these dimensions was significantly correlated with hue (after a Bonferroni correction): the *repeated* *element* (r = .63, p < .001), *diagonal* (r = .45, p < .01) and *upward-opening curve* (r = -.36, p < .05) dimensions. Two of these three dimensions—repeated element and diagonal—were also identified as important predictors in Chapter 3. The third dimension identified in that chapter was the *closed curve* dimension, which does not have a simple correlation with hue distance in this data set. However it is a marginal contributor to the complete 11-predictor model (p = .05), with the only other two contributors to this model being distance along the repeated element (p < .001) and diagonal dimensions (p < .01). Since it had already been identified as a dimension of interest in Chapter 3, it seemed worthwhile to investigate its role as a predictor further.

A 4-predictor model using distance along the repeated element, diagonal, upward-opening curve, and closed curve dimensions (R^2 = .59, p < .001) explained no less variance than the complete 11-predictor model (p > .5). Distance along each of these dimensions was a significant contributor to the 4-predictor model (all ps ≤ .05), and removing any of them from this model results in less variance explained (all ps ≤ .05).

4.4.4 A frequency effect after all

The lack of a frequency-luminance effect came as something of a surprise, given that it is one of the more replicated results in the literature (Beeli, Esslen, & Jäncke, 2007; Cohen Kadosh, Henik, & Walsh, 2007; Simner & Ward, 2008; Smilek, Carriere, Dixon, & Merikle, 2007; Watson, Akins, & Enns, 2012). As there was prior reason to suspect that such an effect exists, I re-ran the frequency-luminance correlation using un-binned data, in order to obtain higher power. This effect was significant (r = .25, p < .001), showing that there is indeed a frequency-luminance effect. It is surprising that the magnitude of this effect increased substantially with unbinned data, as the opposite is normally the case (and is for both the shape-hue and ordinality-hue relationships).

4.5 Discussion

With only one exception, the results from Chapter 3 are entirely replicated in the new data set, with larger effect sizes. We can be very confident that the synaesthetic hue of letters is strongly influenced both by the shape of these letters and by their positions in the alphabet, and that this applies to synaesthetes in general, rather than simply being a peculiarity of those participants in the first study. It is also clear that the strength of these effects are independent of each other, such that a synaesthete whose hues are strongly influenced by letter shape is no more likely than any other synaesthete to have a strong relationship between hue and alphabetical order.

The specific dimensions of shape—repeated element, diagonal and closed curve—identified as driving the shape-hue effect in Chapter 3 are also signficant predictors of hue distance in the present data. Thus it seems likely that there is something about these dimensions that are of particular importance to synaesthetes, but it is not clear what this is. Diagonals and curves are fairly low-level visual features, but the notion of repeated elements (as in the letters *S* and *M*) is not, and other low-level elements of shape such as vertical and horizontal lines do not influence hue. Furthermore, one other shape dimension—upward-opening curve—is an important predictor of hue in the present data but not in Chapter 3. Further research is needed.

The only effect that is not replicated with the present data is the correlation between frequency and luminance, but this appears to be due to two factors: lower power as a result of binning data, and possibly the presence of specific groups of letters that have different relationships with luminance, and thus whose effects might cancel out after binning.

Finally, the synaesthetes in the present data set chose colours that varied in saturation, which I suggested might increase the fidelity of these colours. The larger effect sizes found in this chapter support this suggestion, and also open the door to increasing power by using raw data without binning. The next chapter takes advantage of this in a series of novel analyses.

5 The structure of Czech synaesthesia[4]

5.1 Introduction

We have now established a number of independent relationships between letter similarity and synaesthetic colour similarity among English language speakers. This chapter shifts focus to a different linguistic and cultural context, examining similarity relations between the synaesthetic colours of the 41 letter-colour synaesthetes identified at Charles University in the Czech Republic (see Chapter 2). The Czech language and educational system have a number of unique qualities that enable several important extensions of our understanding of the influences of learning on synaesthetic colour.

5.1.1 Unique features of Czech letters may affect synaesthetic colour

Czech uses an alphabet that is very similar to English (see Table 5.1). Thus the visual stimuli that induce Czech grapheme-colour synaesthesia are more or less the same as in the previous chapters, and indeed the same as in the vast majority of published studies, almost all of which involve languages that use some variation of the Latin alphabet (for some welcome exceptions to this trend, see Asano & Yokosawa, 2011, 2012; Mills et al., 2002; Simner, Hung, & Shillcock, 2011). This similarity of inducers enables testing of many of the same effects presented in Chapters 3 and 4, but the many differences between Czech and English allow us to branch out.

4. See Footnote 1 on p. 26 for authorship details.

Table 5.1 The graphemes of the Czech alphabet including all diacriticals, their alphabetical positions, and order of learning.

Letter	Alphabet Position	Learning Group	Letter	Alphabet Position	Learning Group
A	1	1	*N*	17	3
Á	1	1	*Ň*	18	8
B	2	5	*O*	19	1
C	3	5	*Ó*	19	1
Č	4	5	*P*	20	3
D	5	4	*R*	21	4
Ď	6	8	*Ř*	22	5
E	7	1	*S*	23	2
É	7	1	*Š*	24	5
Ě	7	7	*T*	25	3
F	8	6	*Ť*	26	8
G	9	6	*U*	27	1
H	10	5	*Ú*	27	1
CH	11	5	*Ů*	27	1
I	12	1	*V*	28	4
Í	12	1	*Y*	29	1
J	13	3	*Ý*	29	1
K	14	4	*Z*	30	4
L	15	2	*Ž*	31	5
M	16	2			

5.1.2 Does phonological similarity map to synaesthetic colour similarity?

The first difference to consider between Czech and English is the much higher degree of orthographic transparency of Czech letters. As explained in Chapter 2, in a perfectly transparent language, each grapheme represents only one sound and each sound is produced by only one grapheme (some exceptions to this in Czech are discussed below). This means that the phonological similarity of Czech letters can be measured relatively simply. Doing this is impossible in a language as orthographically opaque as English, where the closest one can come is likely something like the phoneme co-occurence score used in Chapters 3 and 4, which only measures the similarity of letter *names*, not the sounds they represent. Given that various types of letter similarity affect synaesthetic colour, and acoustic properties can also affect synaesthetic colour (e.g. the stressed syllable in an English word often determines its synaesthetic colour, Simner, Glover, & Mowat, 2006), it seems likely that phonological similarity could also affect synaesthetic colours. This might be especially true in an orthographically transparent language like Czech, where phonology and letter identity are almost perfectly mapped to each other.

5.1.3 Alphabetical order vs. learning order in Czech

Another important difference between Czech and English is that Czech students typically learn their letters in a set order that is entirely unrelated to alphabetical position. This enables a clear test of the *learning order* explanation of the ordinality-hue effect. Recall that in Chapter 3, it was suggested that the ordinality-hue effect is the result of the order in which letters are learned. If this is true, then there should be a relationship between Czech learning order and synaesthetic hue. If, on the other hand, there is a relationship between alphabetical order and hue, but no relationship between learning order and hue, this would be very strong evidence against the learning order hypothesis. (Of course finding both effects, or neither, would only confuse matters more.)

The alphabetical order of the Czech letters is presented in Table 5.1. This order may appear haphazard to an English observer, but it has a (relatively) simple phonemic justification. All the vowels have both a short and a long form, where the latter is indicated by the accent known as the *čárka* (which occurs over the first *A* in the word "čárka"). The čárka merely indicates a longer duration of pronunciation, not the underlying vowel quality, and vowels with čárkas are not considered as separate letters in their own right (e.g. they are ignored when determining the order of words in a dictionary or phone book). The small circle over *Ů* (known as the *kroužek*, or "small circle") lengthens the sound of *U* in exactly the same way as a čárka (*Ú* always occurs at the beginning of words, *Ů* always occurs within them), and so is not considered a unique letter either. Czech orthography also contains the *háček* (the diacritical over the *C* in the words "čárka" and "háček"), which can be applied to several consonants and also to *E*. It indicates a palatalized consonant or, to the English speaking ear, a 'softening' of the consonant. The phonemes differ as a result of this palatalization, and consonants with a háček are considered letters in their own right, occupying their own places in the Czech alphabet and dictionary. *Ě* is a special case, as the háček here does not modify the vowel's sound, but rather indicates that the preceding consonant is itself palatalized, and so is not considered as a separate letter from *E*. Finally, the letter pair *CH* (the sound ⟨χ⟩) is also considered as a unique letter.

Table 5.1 also gives the learning order for the Czech letters, adapted from eight recent first-grade textbooks (Březinová, Havel, & Stadlerová, 2007; Ladová, Holas, & Staudková, 2011; Melichárková, Štěpán, & Švecová, 2008; Mikulenková & Mladý, 2004; Mikulenková, Mladý, & Forman, 2004; Nováčková, 2010; Potůčková, 2010; Žáček & Zmatlíková, 2010). As noted above, Czechs learn their letters in a highly regular order. Grade 1 textbooks present the letters in a number of distinct groups, where the letters in each group are presented on the same page or two-page spread of the textbook. While the order of letter presentation within each group varies substantially between texts, the composition and order of the groups themselves is highly consistent (with some exceptions, noted below), and the pedagogical method is highly similar across the Czech Republic. The vowels and their accented forms are always taught first (with the exception of *Ě*), followed by groups of consonants. Each consonant is always introduced as a modifier of the vowels, thus a child learns the letter *T* by learning the nonsense syllables *ta*, *te*, *ti*, *to*, *tu*. By the end of Grade 1 the child is expected to be able to read *any* Czech text fluently, albeit without comprehension, which is of course impossible in a language as orthographically opaque as English.

These learning groups should be taken as a fairly close approximation of any given Czech student's order of learning the letters, rather than a perfect replication of this order. Aside from well-known phenomena such as children learning the letters of their first name prior to other letters (Justice, Pence, Bowles, & Wiggins, 2006), the learning group to which some letters belong differs slightly between textbooks. For example, *Y* is sometimes presented later than the other vowels, and *S* is sometimes a member of group 3, not group 2. Further, the consonants in groups 4 and 5 are not generally presented as a single group on a two-page spread, but rather are each presented on a single page, in an order that varies widely between textbooks (though the letters in group 4 are learned before the letters in group 5). Aside from small differences such as these, however, the learning order is remarkably constant.

5.1.4 A special role for vowels?

Phonological similarity, alphabetical order and learning order apply across all letterrs, but there are also a number of discrete categories among Czech letters that could influence the synaesthetic colours. In each case, the prediction is that letters that are categorized together will be closer in synaesthetic colour.

To begin with, vowels and consonants are clearly distinguished in Czech, as they are in English. But this distinction may be far more salient for Czech than for English speakers, for a number of reasons. There are only five vowel sounds in Czech and at least twice as many in English (the precise number depending on dialect), and thus less phonological variation within the Czech vowels, which might make the vowel class easier to categorize. Furthermore, as described above, consonants are always learned from the start as components of vowel-consonant morphemes, whereas the vowels are taught on their own, and each vowel is re-presented as part of learning each consonant. This special treatment of vowels might also increase the salience of vowels as a special class. Finally, the vowels are the first letters Czech children learn (see Table 5.1), which could serve to further increase the strength of the vowel category.

Vowels being learned first also enables another test of the learning order hypothesis. The previous chapters have shown that letter similarity often leads to synaesthetic colour similarity, which would predict a clustering of the vowels in colour space, since for Czechs they are highly similar in phonology, function within the writing system, and pedagogy. However according to the learning order hypothesis, letters learned earlier are generally further apart in colour space. These two hypotheses, then, make exactly opposite predictions about vowels' colour relationships.

5.1.5 Base-diacritical pairs

Perhaps the most obvious categories in Czech are the ones formed by base letters and their diacritical variations. An intuitive hypothesis is that these would tend to be very close in colour, if not identical, since the base letters and their variations are so similar. An important question, however, is exactly what type of similarity matters here. There are at least three ways in which the base pairs can be similar to their diacritical variations, which do not apply equally to all base-diacritical pairs. These include similarity in terms of shape, identity, and phonology. By choosing comparisons carefully, it may be possible to determine the priority of these three types of similarity in terms of influencing synaesthetic colour.

First, bases and their diacritical forms are highly similar in shape, and we have seen in Chapters 3 and 4 that shape similarity is generally correlated with hue similarity. Shape similarity then, would seem to apply more or less equally to all base-diacritical pairs. Second, as discussed above, vowels with čárkas are considered to be variants of the same letter as their base forms, whereas consonants with háčeks are distinct letters (see Table 5.1). This, then, might cause vowel-čárka pairs to be closer together in colour than consonant-háček pairs. Third, there is a phonological relationship between bases and their diacriticals, but the precise nature of this relationship differs between čárkas and háčeks. Vowels, as noted above, are merely lengthened by their čárkas, with no change in vowel quality, whereas háčeks indicate a different place of articulation for the phoneme. This change is rule-governed (palatalization), but it is more significant than the lengthening of vowels by the čárkas. Thus if phonological similarity drives synaesthetic colour similarity, we might again expect vowel-čárka pairs to be closer together in colour than consonant-háček pairs.

Merely comparing vowel-čárka and consonant-háček pairs, then, confounds similarity in terms of abstract identity and in terms of phonology, since vowel-čárka pairs are more similar in both ways. However a third type of category within Czech letters might allow these two types of similarity to be disentangled. The members of each pair of voiced-unvoiced consonants (*D-T*, *G-C*, *H-CH*, *V-F*, and *Z-S*) bear very little shape similarity to each other, and are not considered to have the same letter identity, but their phonological relationship is quite similar to that between consonants and their háčeked forms. Thus if phonological relationships influence letter similarity, one would expect to find that pairs of voiced-unvoiced consonants are somewhat closer to each other than they are to the other consonants. Thus by comparing how clustered vowel-čárka pairs are from the other vowels, and how clustered consonant-háček and voiced-unvoiced pairs are from the other consonants, one can determine which of the three types of similarity affect synaesthetic colour. Table 5.2 summarizes the possibilities.

Table 5.2 Three types of letter pairs in Czech, three ways in which the letters in these pairs are similar to each other, and three predictions about what ought to happen if each type of similarity affects synaesthetic colour.

	Similar shapes	Same letter identity	Similar phonology
Vowel-čárka pairs (*A-Á*, *E-É*, *I-Í*, *O-Ó*, *U-Ú*, *Y-Ý*)	Yes	Yes	Identical phoneme
Consonant-háček pairs (*C-Č*, *D-Ď*, *N-Ň*, *R-Ř*, *S-Š*, *T-Ť*, *Z-Ž*)	Yes	No	Different phoneme
Voiced-unvoiced pairs (*D-T*, *G-C*, *H-CH*, *V-F*, *Z-S*)	No	No	Different phoneme
Prediction for synaesthetic colour (assuming similarity maps to similarity).	**Vowel-čárka and consonant-háček pairs should cluster, voiced-unvoiced pairs should not.**	**Vowel-čárka pairs should cluster, consonant-háček pairs and voiced-unvoiced pairs should not.**	**All three types of pairs should cluster, vowel-čárka pairs closest of all.**

5.1.6 I and Y

One final category that may influence Czech synaesthetic colour stems from one of the rare exceptions to the orthographic transparency of Czech. *I* and *Y* both represent the phoneme ⟨ɪ⟩, and *Í* and *Ý* both represent the phoneme ⟨iː⟩. This phonemic identity, then, might translate into colour similarity, with *I* and *Y* (and their čárka forms) closer to each other than they are to the other vowels.

5.1.7 Outline of this study

The remainder of this chapter explores each potential influence on Czech synaesthetic colour in turn. First, I look for evidence of each of the potential categorical effects just described. I then return to the correlational analyses that were the centerpieces of Chapters 3 and 4, incorporating novel measures of letter similarity derived from Czech letters' frequency, alphabetical order, learning order, and phonology.

5.2 Participants

41 native Czech-speaking letter-colour synaesthetes were identified as part of the Charles University survey described in Chapter 2. As with the English-speaking participants in Chapter 4, all participants were confirmed as having highly-consistent grapheme colours by the Synesthesia Battery (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007).

5.3 Data preparation

As in Chapter 4, the final (third) colour assigned to each letter on the Synesthesia Battery was used in all analyses, and if a participant chose "no colour" for any trial, that letter was removed from all analyses for that participant. Once again, participants' letter colours were transformed from RGB to CieLAB coordinates, and three separate colour distance measures were computed: Colour, Luminance, and Hue distance. Mean distance scores for each letter pair were computed by averaging across all participants. As in Chapter 4, participants had the option of selecting "no colour" for particular letters, in which case their data was ignored for that letter.

Five letter similarity measures were computed, summarized in Table 5.3. The *Shape Difference* measure was adapted from the one used in Chapters 3 and 4 (Gibson, 1969), with the addition of three extra dimensions: one representing the presence or absence of a čárka (´), one the presence or absence of a háček (ˇ), and one the presence or absence of the kroužek (˚). Also as in Chapters 3 and 4, *Frequency* and *Ordinality Difference* scores were computed as Weber fractions (absolute value of the difference between two letters divided by the sum). Data for the frequency fractions were taken from the overall frequency of each grapheme in a large corpus of Czech texts, as reported by Králik (1983). Ordinality was computed using Czech dictionary order, as given in Table 5.1.

Learning Order Difference was another Weber fraction, computed using the learning groups presented in Table 5.1. Finally, *Phonological Similarity* of Czech letter pairs was computed using the SimilarityCalculator PERL script (Albright, 2006). This characterizes similarity according to the system of Frisch (Frisch, 1996; as used in, e.g., Frisch, Pierrehumbert, & Broe, 2004), which defines the similarity of two phonological segments as the number of natural classes they share over the sum of all shared and non-shared natural classes. A complete set of feature values for the Czech letters was custom-developed for this study by a collaborator (John Alderete, SFU) using the Unified Feature Theory of Clements & Hume (1995), which specifies each phonological segment in terms of 19 articulatory features.

Table 5.3 Letter similarity measures used in the Czech study.

Similarity Measure	Description
Shape difference	Euclidean distance in an 14-dimensional space defined using the basic letter shape features from Gibson (1969) shown in Table 3.2, with the addition of čárka (´), háček (ˇ), and kroužek (˚) features.
Frequency difference	Difference of two graphemes' frequencies in Czech divided by the sum of their frequencies (from Králik, 1983).
Ordinality difference	Difference of two graphemes' positions in the alphabet divided by the sum of their positions (see Table 5.1).
Learning order difference	Difference in two graphemes' learning order group in Czech, divided by the sum of these groups (see Table 5.1).
Phonological similarity	Phonological similarity between Czech letters.

Participants provided colours for 38 letters: all 23 base letters plus all their diacritical forms (not including *CH* due to experimenter error). There are 703 possible letter pairs made from these 38 letters, and thus each similarity measure has 703 points. Unlike Chapters 3 and 4, no binning was performed on the data. This allows for increased power, but at the expense of effect sizes, which should be kept in mind when making comparisons to previous chapters.

5.4 Results

5.4.1 Categorical second-order influences on synaesthetic colour in Czech

Our Czech synaesthetes formed several categorical clusters of graphemes in colour space, such that members of a given category were closer together than they were to other graphemes (see Figure 5.1). This is supported by a series of t-tests comparing the mean distances of the categories noted in the Introduction. As is immediately obvious upon glancing at Figure 5.1, there was a bimodal distribution of letter pair distances in colour space, such that bases and their diacritical variations were much closer to each other than to other letters, in terms of both hue (t_{656} = -16.33, p < .001) and luminance (t_{656} = -14.61, p < .001). The base-diacritical pairs could be further sub-divided: vowels with čárkas were closer together than consonants with háčeks (Hue: t_{11} = -3.42, p < .01, Luminance: t_{11} = -3.60, p < .01).

The bimodal nature of these data meant that base-diacritical pairs had to be removed from all analyses not specifically pertaining to them, because any other effects could easily be swamped by the strong clustering of base-diacritical pairs. This is particularly true of the correlational analyses in the next sections, since base-diacritical pairs were also very close in terms of shape, ordinality, learning order, and phonology, which constitute most of the measures to be correlated with colour. Thus the 17 base-diacritical pairs were removed from the remainder of analyses, leaving 686 grapheme pairs in total.

Among these 686 pairs, there was a less visually obvious, but equally significant, clustering of vowels, which were closer to each other than they were to the consonants (Hue: t_{415} = -3.34, p < .001, Luminance: t_{415} = -4.60, p < .001). Conversely, the consonants tended to be slightly further apart in hue from each than they were from the vowels (t_{603} = 3.20, p < .01), but were not clustered in luminance (p > .4).

Breaking things down further, there was at least one sub-cluster within the vowels: I and Y (and their accented forms) are closer together than they were to the other vowels in hue (t_{79} = -3.68, $p < .001$), but very slightly further apart in luminance (t_{79} = 2.54, $p < .05$). The only potential letter clusters mentioned in the Introduction that do not appear to impact synaesthetic colour were pairs of voiced-unvoiced consonants, which were no closer together in either hue or luminance than they were to other consonants (both *ps* > .25, not shown in Figure 5.1).

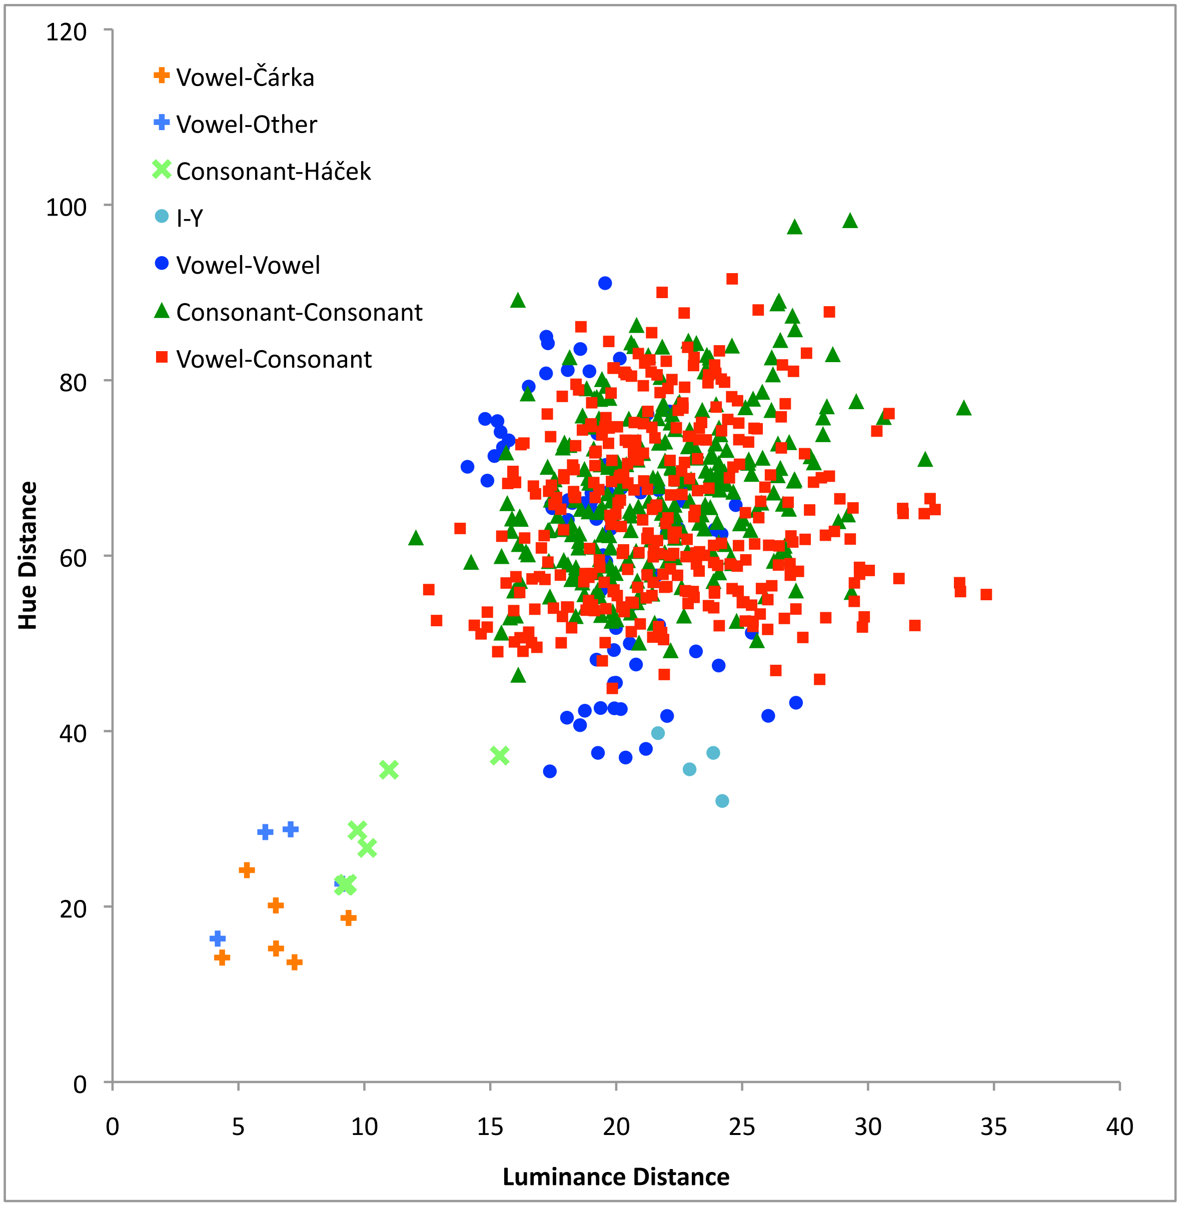

Figure 5.1 Mean hue and luminance distances of all Czech letter pairs across all participants, grouped into various categories. The "Vowel-Other" pairs include E-Ě, É-Ě, U-Ů,and Ú-Ů. The I-Y pairs are also included in analyses of the Vowel-Vowel group.

With one possible exception, English synaesthetes did not appear to cluster any of these categories in colour space. While none of the tests involving base pairs with diacriticals could be performed for the English alphabet, those tests that were meaningful in English showed no significant effects in the colour data used in Chapter 4 save, curiously enough, for consonants being slightly closer together in luminance than they were to vowels ($t_{308} = -2.80$, $p < .01$; all other ps > .25).

5.4.2 Notes on the remaining analyses

As described above, the 17 base-diacritical pairs were removed from all remaining analyses. However preliminary analyses revealed that this was not always enough, and that some important effects could not be uncovered without removing diacriticals from analyses entirely. Consider that among the 686 grapheme pairs there are twice as many pairs involving some version of *A* as pairs involving some version of *B*, since *A* can be modified by the čárka, but there is no version of *B* with a háček. *E* and *U* would be represented three times as often as a letter like *B*, since they each have two diacritical variations. Since diacritical variations tend to be very similar in colour to their base forms (see Figure 5.1), including them in the data set effectively magnifies the importance of the colours of letters that have diacritical versions, giving them two or three times as much potential influence over the results as their unmodified cousins. One way of proceeding that could avoid this would be to remove diacriticals entirely from all analyses, leaving only the 23 base letters. This would be far from ideal, however, both because it greatly reduces power and because the diacriticals are interesting in their own right. Thus, the remainder of analyses in this chapter were performed twice, once on the complete data set of 686 grapheme pairs (not including the 17 base-diacritical pairs), and once on the 253 pairs of base letters only. For the sake of brevity, however, results will only be reported for the complete set of 686 pairs, except in those instances where there are important differences. Where the smaller data set is not explicitly mentioned, the reader can safely assume that the qualitative pattern of results was the same as in the larger set (i.e. all significant effects in either data set were significant, or at least marginal, in the other, and in the same direction).

Another wrinkle that became apparent during preliminary analyses is that a number of effects were different for vowel-vowel, consonant-consonant, and vowel-consonant pairs. Thus the in-depth analyses of each important relationship between letter and colour similarity split the letter pairs into these three groups, although it should be kept in mind that this reduced power once more. The lack of an effect for one (or all three) groups does not necessarily mean that they did not contribute to an effect that applies over the entire set of letter pairs. However when strong differences between the groups were found, these may be indicative of some of the important processes that underlie Czech synaesthetic development.

5.4.3 Shape-, ordinality-, and phonology-colour correlations in Czech

As in Chapters 3 and 4, simple correlations were calculated between the various letter similarity measures and colour distance measures, multiplying *p*-values by 15 to compensate for the number of tests. For the simple correlations, the pattern of results was very similar for both the complete data set including diacriticals and the reduced base letter only data set.(see Table 5.4). As one would expect, correlations in the smaller base letter data set were generally less significant, but were all in the same general range as the results from the larger set. In both sets, shape distance was correlated with colour distance, and this was solely due to a correlation with hue distance, as in Chapters 3 and 4. Both ordinality difference and phonological similarity were correlated with colour distance in the larger data set, but these correlations arose from relationships with both luminance and hue. In the base letter set, none of the correlations with ordinality difference were significant, although both the ordinality-luminance and ordinality-hue correlations were marginally significant. Furthermore, the correlation between phonological similarity and luminance distance was no longer significant, although given that its absolute magnitude increased from 0.11 to 0.14, this was likely a power issue (without the Bonferroni correction, it was significant: $p = .03$).

Table 5.4 Simple correlations between Czech letter similarity measures and colour distance, calculated using all 38 graphemes (after removing the base-diacritical pairs) and the 23 base letters only.

Similarity Measure	All Graphemes (no base-diacriticals) (38 graphemes, 683 pairs)			Base Letters Only (23 graphemes, 253 pairs)		
	Colour Distance	Luminance Distance	Hue Distance	Colour Distance	Luminance Distance	Hue Distance
Shape difference	0.21***	0.03	0.21***	0.19*	0.07	0.19*
Frequency difference	-0.07	0.02	-0.09	0.09	0.04	0.07
Ordinality difference	0.16***	-0.19***	0.21***	0.12	-0.16 .	0.16 .
Learning order difference	0.05	0.10	0.01	-0.02	0.06	-0.04
Phonological similarity	-0.15***	-0.11*	-0.14**	-0.21**	-0.14	-0.20**

All p-values are Bonferroni corrected and > .1, except: . $p < .1$, * $p < .05$, ** $p < .01$, *** $p < .001$.

5.4.4 Independent influences of shape, ordinality, and phonology

The various effects reported in Table 5.4 were independent of each other. That is, each of the aspects of letter similarity that are significantly (or marginally) correlated with synaesthetic luminance or hue accounted for different portions of the variance in luminance or hue distance. This is supported by two linear models, one which predicted hue distance on the basis of shape difference, ordinality difference and phonological similarity, and the other which predicted luminance distance on the basis of ordinality difference and phonological similarity (Table 5.5). All the letter similarity measures were significant contributors to their models.

Table 5.5 Summary of linear models predicting hue and luminance, giving t-values of each predictor and R^2 of the models.

Difference Measure	Luminance Model t-values	Hue Model t-values
Shape Difference	N/A	6.04***
Ordinality Difference	-4.74***	6.26***
Phonological Difference	-2.47*	-4.81***
Model R^2	0.04***	0.12***

All p-values are > .1, except: . $p < .1$, * $p < .05$, ** $p < .01$, *** $p < .001$.

In terms of absolute effect sizes, both models were fairly anemic. Keep in mind, however, that effect sizes would be larger with binned data as is used in Chapters 3 and 4. More importantly, these models establish that the effects were independent of each other, and so a complete account of influences on synaesthetic colour in Czech needs to consider each of these three letter similarity measures separately. The following three sections consider each of these effects in turn.

5.4.5 Shape-hue effect is strongest for vowel-vowel pairs, predictive dimensions for Czech and English overlap

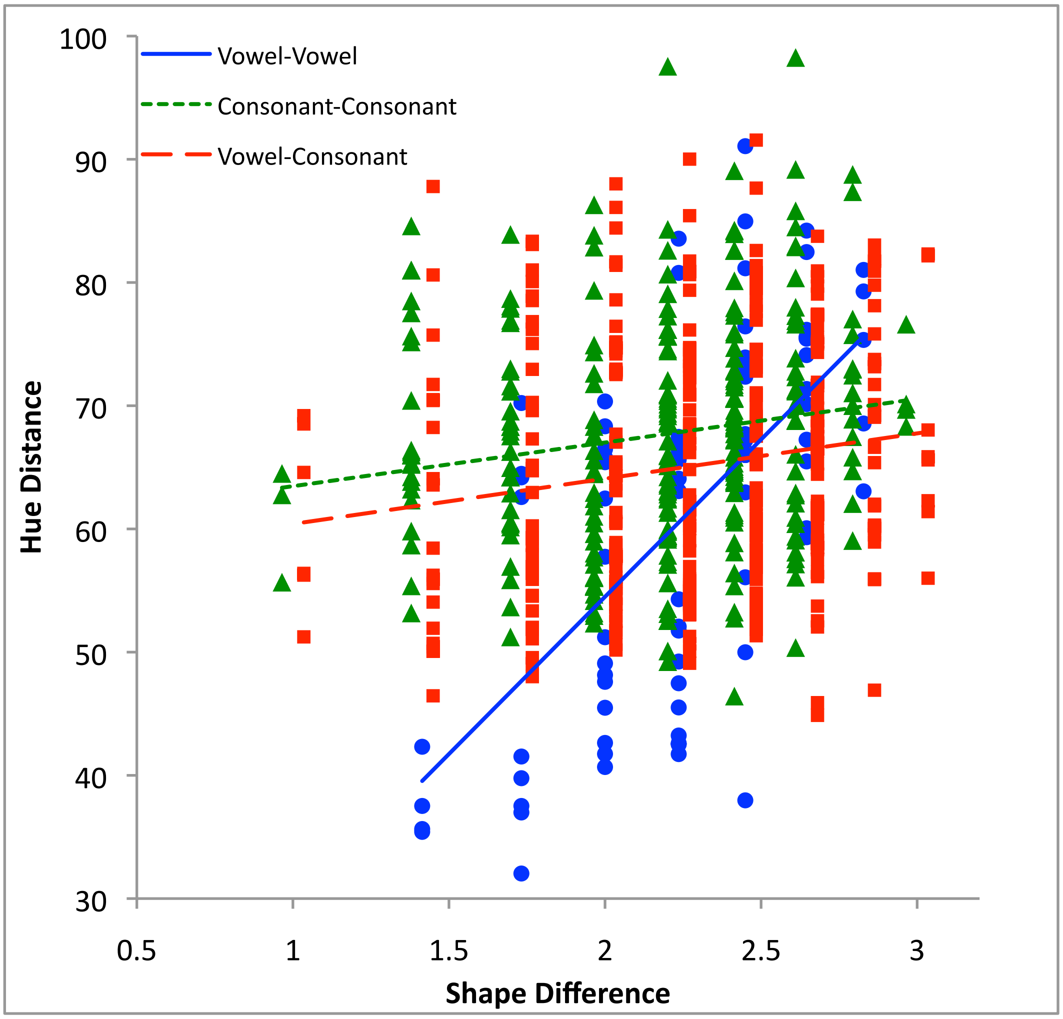

Figure 5.2 The shape-hue effect for Czech letters (not including base-diacritical pairs). In order to improve legibility, the Shape Difference values for Consonant-Consonant pairs have been shifted slightly to the left, and the values for Vowel-Consonant pairs have been shifted slightly to the right.

The shape-hue effect was due to a general trend across all letter pairs, rather than to a few outliers (see Figure 5.2). The effect was much stronger among vowel-vowel pairs (*r*

= .63, $p < .001$), but was still significant for consonant-consonant pairs ($r = .16$, $p < .01$) and for vowel-consonant pairs ($r = .16$, $p < .01$).

While the shape-hue effect had virtually the same strength for both data sets (see Table 5.4), the two differed markedly in terms of their usefulness for determining the dimensions of shape that actually drove the overall effect. Indeed, this could only be done using the smaller data set that excluded the diacriticals. This is supported by a series of linear models, described below.

The data set using all 38 graphemes did not allow for a sensible interpretation of the dimensions of shape that influence synaesthetic hue. Distance along all but 4 (*Straight Vertical*, *Closed Curve*, *Diagonal*, and *Háček*) of the 14 shape dimensions had a significant simple correlation with hue distance, even after a Bonferroni correction (all *p*s < .01). And all but 5 (*Straight Horizontal*, *Straight Vertical*, *Symmetry*, *Diagonal*, and *Háček*) were significant contributors to a model predicting hue distance on the basis of all the shape dimensions ($R^2 = .29$, $p < .001$). Without compelling theoretical reasons to select from among the 9 or 10 dimensions that were significantly associated with Czech synaesthetic hue, it was impossible to reduce this models further, and one can be reasonably certain that some of the variance it explained was simply due to the sheer number of predictors in play. Note that the Diagonal and Close Curve dimensions, which were among the few non-significant predictors of shape here, were important predictors in both sets of English data from Chapters 3 and 4.

However the smaller data set including only the base letters produced more interpretable results. Bonferroni-corrected simple correlations between distance along the various shape dimensions and hue distance revealed only three signficant correlations: *upward-opening curve* ($r = -.20$, $p < .05$), *repeated element* ($r = .29$, $p < .001$), and *horizontal discontinuity* ($r = .24$, $p < .01$). These three dimensions were also the only signficant predictors (*p*s < .01, all other *p*s > .05) in a model predicting hue using all 11 shape dimensions ($R^2 = .20$, $p < .001$). (Note that this model has 11 dimensions instead of 14, as the three diacritical dimensions had no predictive value for the base letters.) A model using only these three dimensions to predict hue distance (R^2 = .16, p < .001) explained only 3% less variance than the complete 11-dimensional model (see Table 4.1), but this difference was nevertheless marginally significant (p = .08). Using these three dimensions alone as predictors in the complete data set explained almost as much variance as the complete 14-predictor model (R^2 = .24, p < .001), but the remaining 5% of variance was a significant difference between the 3-predictor and 14-predictor models (p < .001), again indicating that some other dimension of letter shape was an important predictor. Thus these three dimensions explain *almost* all the variance in hue space attributable to shape similarity for Czechs. The repeated element dimension was a contributor to the reduced shape models for the English data in Chapters 3 and 4, while the upward-opening curve dimension was a contributor in Chapter 4, indicating a partial overlap between the dimensions of shape that influence synaesthetic colour in all three samples.

For the sake of completeness, an attempt was made to account for the remaining difference between the reduced model and the complete 11-predictor one. Three dimensions were marginal contributors to the 11-predictor model (*diagonal*, *horizontal-opening curve*, and *horizontal discontinuity*), and adding any one of these to the 3-predictor model explained no less variance than the complete model (all ps > .2), while adding any other dimension to the 3-predictor model resulted in a model that was still marginally different from the complete model (.05 < p < .1 in all cases). While the diagonal dimension was important in both sets of English data, there was no principled reason to prefer it over the other two marginal predictors of hue distance, and so the present data did not support any further conclusions.

5.4.6 Ordinality effects with both hue and luminance, in opposite directions

The ordinality-hue effect, like the shape-hue effect, was strongest for vowel-vowel pairs (r = .63, p < .001), and weaker for vowel-consonant pairs (r = .32, p < .001) (see Figure 5.3). Unlike the shape-hue effect, consonant-consonant pairs made no contribution to the overall ordinality-hue effect (p > .3), indeed they trended in the opposite direction.

The ordinality-luminance effect, on the other hand, was only significant for vowel-consonant pairs (r = -.25, p < .001, other ps > .3) although the trend for vowel-vowel and consonant-consonant pairs was still negative.

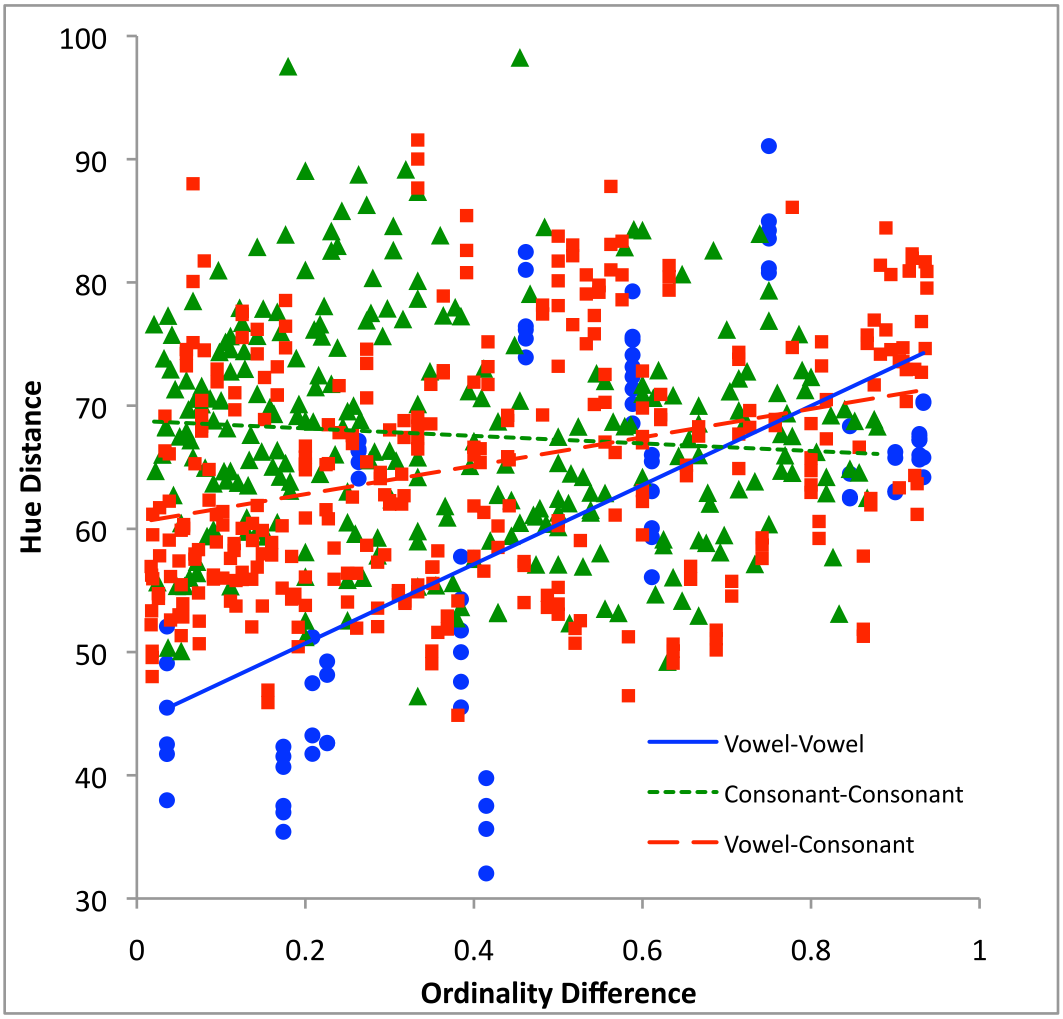

Figure 5.3 The ordinality-hue effect in Czech synaesthesia.

No further reductions of the ordinality effects are possible at present. It seemed that the ordinality-luminance effect might be driven solely by the unusual brightness of *I* and *O*, which have often been reported as preponderantly white (cf. Simner et al., 2005), which would lead to a large luminance distance with other letters. In the Czech data, *I*, *Í*, *O* and *Ó* were all among the 5 brightest letters, with a mean CIE L value of 85, compared to an average of 71 for all other letters (t_{37} = 3.84, p < .001). However after removing them from the data set, the ordinality-luminance effect was still significant for vowel-consonant pairs ($r = -.19$, $p < .01$). No other potential outliers were apparent in these data.

5.4.7 Phonology-colour relations are strongest for vowel-like consonants

The correlation between phonological similarity and hue was significant for both the complete and the base-letter-only data sets, but when these were decomposed along vowel-consonant lines only one significant correlation was found, for vowel-consonant pairs in the base-letter-only data set ($r = -.26$, $p < .01$). The phonology-luminance correlation disappeared altogether when the data were decomposed into the vowel/consonant groups (all *p*s > .05).

Of course this does not mean that the phonology-hue and phonology-luminance effects were spurious, simply that these effects were likely weak general trends across all letter pairs, and carving the data into three groups lowered power. However carving up the data in this way still proved valuable, as further investigation of the one significant correlation revealed a very clean effect of phonological similarity between vowels and consonants. In Figure 5.4, the vowel-consonant pairs are presented by themselves, and further split into 3 sub-groups: all pairs where either the consonant or the vowel has a diacritical, all pairs involving a consonant other than *J* or *V*, and all pairs involving *J* or *V*. This scatterplot allows us to see two trends clearly. First, among the base pairs without diacriticals, the phonology-hue effect was driven almost entirely by *J* and *V*, which are the closest base consonants in phonology to all the vowels. These two letters appeared to have been pulled towards the vowels by their relatively high phonological similarity, leading to the strong phonology-hue effect among the vowel-consonant pairs. Removing pairs involving *J* or *V* from the base-letter-only data set rendered the phonology-hue effect among vowel-consonant pairs insignificant ($r = -.14$, $p = .19$), although there was still a significant effect across all base letter pairs ($r = -.13$, $p < .05$).

Second, there was no phonology-hue effect for the vowel-consonant pairs if the pairs including diacriticals are included, despite the fact that several diacriticals had a far higher degree of phonological similarity to the vowels than either *J* or *V*. Presumably, this was because the categorical similarity between bases and diacriticals (see Figure 5.1) overrode the phonological similarity between these vowel-consonant pairs involving diacriticals. This would explain why the absolute magnitude of the correlation between phonological similarity and colour distance went up when the diacriticals were removed from the data set (see Table 5.4).

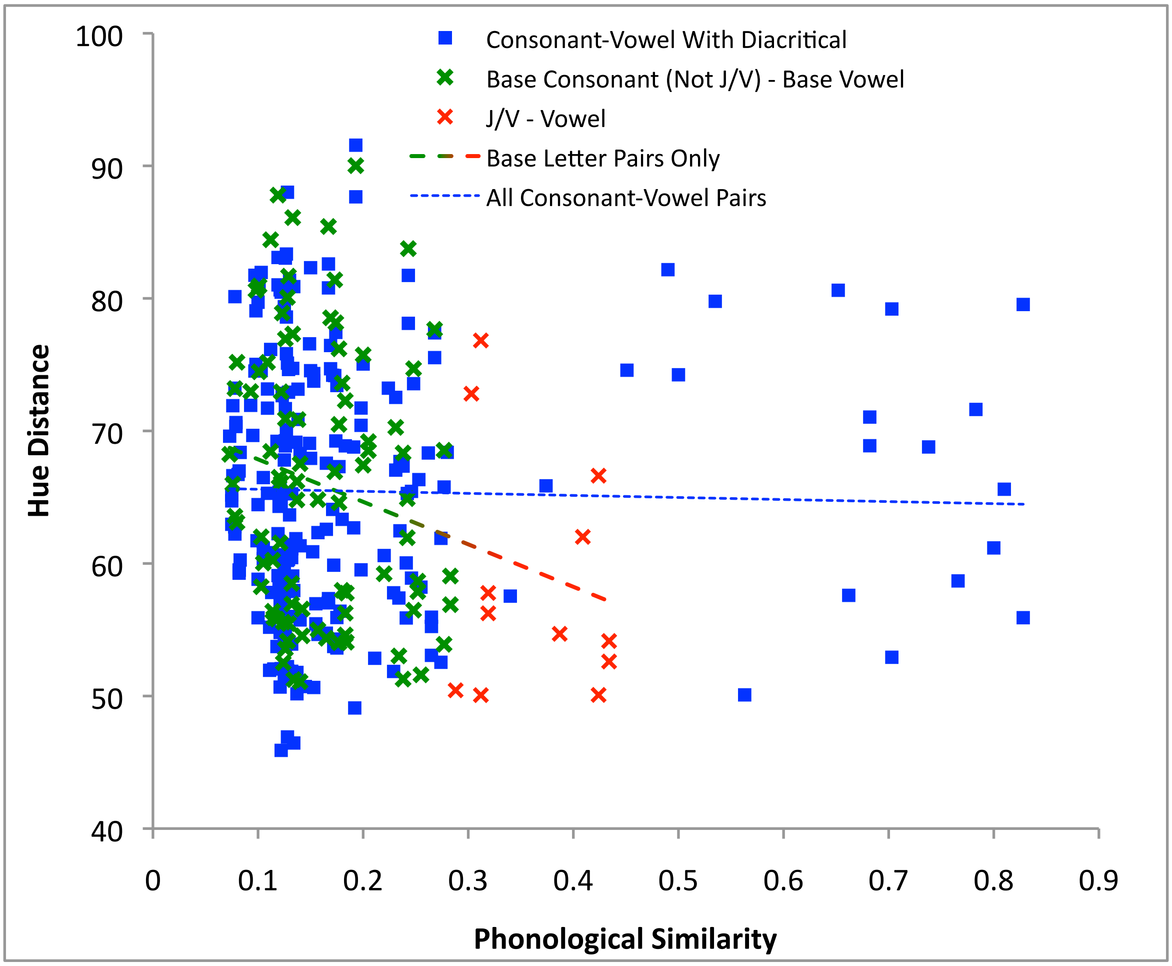

Figure 5.4 The phonology-hue effect for vowel-consonant pairs alone, demonstrating the special role of *J* and *V*, and the lack of an effect among pairs that include diacriticals.

Further investigations did not uncover any other outliers or unusual patterns in either phonology-hue or phonology-luminance relationships.

5.4.8 Is there any impact of learning order on synaesthetic colour?

The learning order hypothesis was not well-supported by the data thus far. The vowels, which are learned first by Czechs, were closer together in both hue and luminance than they were to the subsequently-learned consonants, exactly opposite to the prediction of the hypothesis. Furthermore, there was no correlation between learning order and hue or luminance, while there was also a moderate correlation between alphabetical order and hue, despite the fact that Czechs do not learn their letters in alphabetical order.

In order to leave no stone unturned, the correlational analyses between learning order difference and hue and luminance distance were re-run, splitting the data set into the three consonant/vowel groups. Surprisingly, there were highly significant effects. There was no relationship with either hue or luminance among the consonant-consonant pairs (all *p*s > .1), nor among the consonant-vowel pairs when the base letter only data set was used (both *p*s > .5). However when the diacriticals were included, there was a relationship between learning order distance and hue distance among the vowel-consonant pairs ($p < .05$, $r = -.13$) the vowel-vowel pairs ($p < .001$, $r = .40$), as well as a relationship between learning order distance and luminance distance among the vowel-vowel pairs ($p < .05$, $r = -.26$; $p > .5$ for vowel-consonant pairs).

It quickly became apparent that all the apparent relationships here were driven solely by *Ě*, the only vowel that is not learned at the beginning of Czech letter-learning (see Table 5.1). Removing pairs including *Ě* from the data set eliminated all effects (all *p*s > .1). The reason for this is easy to understand in the case of vowel-consonant pairs. Because vowels, including *Ě*, clustered together in hue and luminance (see Figure 5.1), then *Ě* would tend to be far away in hue and luminance from the consonants that it was close to in learning order, producing the negative relationship between learning order difference and hue distance among the vowel-consonant pairs. With the vowel-vowel pairs, however, the role of *Ě* was due to two effects we did not originally anticipate: *Ě* was further in hue from the other vowels, on average, than the other vowels with diacriticals were from these vowels, and closer in luminance (hue: $p < .001$, $t_{64} = -3.99$, luminance: $p < .05$, $t_{64} = 2.29$). Note that no form of *p*-value correction was used here, and it is quite possible that either of these effects is spurious. Certainly there is a sensible explanation for the hue effect: *Ě* is categorically distinct from vowels with čárkas or kroužeks, in that it indicates a palatalization of the immediately preceding consonant, as opposed to merely lengthening the vowel. Thus it may be that its separation in hue was due to another categorical effect like those shown in Figure 5.1. However this does not explain the fact that it is closer together in luminance. Whatever the reason for the effects involving *Ě*'s, these analyses give no reason to think that there is any genuine impact of learning order on synaesthetic colour in Czech.

5.4.9 Chasing down the frequency-luminance effect

The lack of a frequency-luminance effect (see Table 5.4) once again came as a surprise, although this time it could not be blamed on low power alone, since the data are not binned as they were in Chapters 3 and 4. One further analysis did indicate that such an effect existed, but only as a first-order effect, unlike all the other effects described in this chapter. This was by calculating, for each subject, the magnitude of the correlation between the raw luminance values of each letter (not distances between letter pairs) and the raw frequency of each letter. Only one of these correlations was significant on its own ($p < .05$ with no Bonferroni correction), but on average, the absolute magnitudes of the individual correlations were above 0 ($p < .01$, $t_{40} = 3.17$). This was not the case when individual correlations were calculated between frequency difference and luminance distance ($p = .20$), nor when correlations were calculated with either dimension of hue (both *p*s > .1). Thus there was a very small frequency-luminance effect in these data, but only at a first-order level.

5.5 Discussion

5.5.1 Overview of results

Czech grapheme-colour syanesthetes, like their English counterparts, are influenced by a wide range of learned letter properties as they develop their letter-colour associations. To begin with, they cluster their letters categorically within colour space. Base letters are extremely close to their diacritical variations, and among these pairs, vowel-čárka pairs are closer to each other than consonant-háček pairs. Vowels are slightly closer to each other than they are to consonants, and within the vowels, *I*, *Y* and their accented forms are closer to each other than they are to the other vowels. As with English synaesthetes, they tend to associate similarly-shaped letters with similar hues, and less similar colours to letters that are earlier in the alphabet. Furthermore, similar-sounding letters tend to be more similarly-coloured, both in luminance and hue. These three effects all differ for different groups of letters: the shape-hue and ordinality-hue relationships are particularly strong for vowel-vowel pairs, and the phonology-hue and phonology-luminance relationship are much stronger for vowel-consonant pairs, specifically those involving the consonants that are most vowel-like. Unlike English synaesthetes, there is no second-order relationship between letter frequency and synaesthetic luminance, though there is a small first-order effect. Finally, there is no hint of a relationship between the order in which letters are learned and their synaesthetic colours.

5.5.2 No learning order effect, a special role for sequences?

These results force the rejection of the learning order hypothesis, which states that synaesthetes tend to assign more distinct colours to letters earlier in the alphabet because this is (roughly) the order they learn these letters in. Czechs do not learn their letters in alphabetical order, but the learning order of Czech letters is not even close to being correlated with synaesthetic colour, while alphabetical order is. Furthermore, Czech vowels, which are the first letters Czech children learn, are clustered together in both hue and luminance, while the learning order hypothesis suggests that they should be driven widely apart. One can account for the ordinality-hue effect among both Czech and English speakers by positing a single mechanism related to ordinality itself, rather than learning order. That is, ordinality maps on to synaesthetic colour because ordinality is particularly salient to synaesthetes. This is consistent with the position argued by Eagleman and colleagues in recent years, who have suggested that grapheme-colour synaesthesia is a sub-type of “coloured sequence synaesthesia” (Novich, Cheng, & Eagleman, 2011; Pariyadath, Plitt, Churchill, & Eagleman, 2012; Tomson et al., 2011).

As the name implies, coloured sequence synaesthesias are those in which individuals associate colours with items that are habitually learned as members of a sequence. This may be the most common type of synaesthesia, and appears to be only weakly related to forms involving colour but not sequences, such as pain-colour or orgasm-colour, or forms involving sequences but not colour, such as spatial forms for numbers or time (Novich, Cheng, & Eagleman, 2011). Neuroimaging work in non-synaesthetes suggest that over-learned sequences, unlike linguistic items in general, are predominantly processed in the right hemisphere, specifically in the middle temporal gyrus and inferior parietal lobe (Pariyadath, Plitt, Churchill, & Eagleman, 2012). Several neuroimaging studies have shown these areas activated in synaesthetic perception involving sequential stimuli, however several others have not (for a review see Rouw, Scholte, & Colizoli, 2011). These Czech data can be seen as further support for the idea that there is something special about sequences for synaesthetes.

5.5.3 Ranking the influences on synaesthetic colour

These results also allow a tentative ranking of some of the influences on synaesthetic colour. For example, it appears that phonological similarity is of less importance than visual or ordinal similarity. Two results point towards this conclusion. First, consonant-háček pairs are tightly clustered in colour space, but pairs of voiced-unvoiced consonants are not. Both groups have similar phonological relationships within each letter pair, but only the consonant-háček pairs are similar in shape and ordinality, suggesting that one or both of these two types of similarity accounts for the consonant-háček effect. Second, there is a relatively strong phonology-colour relationship for vowel-consonant pairs, due to the pair of vowel-like consonants *J* and *V* being “pulled” towards the vowel cluster. However this effect vanishes when letters with diacriticals are included in the analysis. Despite the fact that many of these letters are much more vowel-like than either *J* or *V* (Figure 5.4), there is no hint that this affects their synaesthetic colours. I suggest this is because their colour is far more determined by their strong relationship (in terms of shape, ordinality, or abstract identity) with their base letters. While phonology clearly affects synaesthetic colour, it is trumped by shape and ordinality in the present data.

This is in keeping with a report of a native English-speaking synaesthete who learned Russian in high school, long after her English letter colours had stabilized. Her colours for Cyrillic letters were strongly influenced by their visual and phonological similarity to English letters, however in a case where a Cyrillic letter was visually similar to one English letter and phonologically similar to another, it usually took the colour of the similarly-shaped letter (Mills et al., 2002).

The relative importance of shape and ordinality cannot be determined from these data. The shape-colour effect in the present data is slightly stronger in absolute magnitude, but the ordinality-hue effect is stronger for the English synaesthetes in Chapter 4. Vowel-čárka pairs are closer in hue than consonant-háček pairs, and they are closer in terms of ordinality while being (arguably) no closer in terms of shape. However this does not show that ordinality is in any way overriding the effects of shape, simply that it can add to it. These data, then, do not allow the shape and ordinality to be disentangled.

5.5.4 A special role for vowels?

There seems to be a special role for vowels in Czech synaesthesia. Vowels are somewhat clustered together in colour space, but they also seem to be the prime movers of several of the other similarity relationships. For instance, the shape-hue and ordinality-hue relationships are far stronger for vowel-vowel pairs than they are for consonant-consonant pairs, indeed of the effects discussed in this chapter, only the shape-hue effect has a detectable impact on consonant-consonant pairs at all. Furthermore, the phonology-hue effect appears to be driven largely by the similarity of *J* and *V* to the vowels.

5.5.5 Ambiguous evidence for a hue/luminance split

Evidence of a hue/luminance split is more ambiguous in the Czech data than those from the English synaesthetes in Chapters 3 and 4. As in the English data, there is a shape-hue effect but no shape-luminance effect, and a frequency-luminance but no frequency-hue effect. Further, while there are both ordinality-hue and ordinality-luminance effects of roughly equal magnitude, they are in opposite directions, indicating that they cannot be due to the same factors. Thus for the similarity dimensions that were tested among

English synaesthetes, we replicate the finding of a hue/luminance split. However, all the categorical effects shown in Figure 5.1 are more or less equivalent in hue and luminance, save that *I* and *Y* are only clustered within hue space. Also, the phonology-hue and phonology-luminance effects seem essentially indistinguishable from each other, especially among the base letter pairs where they are strongest. Further work is needed to determine why the hue/luminance split occurs for some types of similarity but not others.

What these results establish beyond the shadow of a doubt is that synaesthetic colours of Czech speakers reflect a number of learned properties of letters. Thus Czech grapheme-colour synaesthesia, like English grapheme-colour synaesthesia, encodes a surprising amount of information about its inducer domain.

6 Grapheme-colour synaesthesia benefits rule-based category learning[5]

6.1 Introduction

We have now seen strong evidence for the effect of learning on synaesthesia. It has been established that the likelihood of developing (or retaining) synaesthesia is dependent upon childhood learning challenges (Chapter 2), and that specific synaesthetic associations encode a great variety of information about the inducer domain (Chapters 3-5). Now we examine the other direction of the relationship, verifying if synaesthesia can be useful for learning.

Is *synaesthesia* good for anything? The suggestion that it has some utility goes back well over a century (Calkins, 1893; Calkins, 1895) and recent work has begun to confirm this. *Grapheme-colour synaesthetes*, who experience letters and numerals as having specific colours, have episodic memory advantages for letters and words (Radvansky, Gibson, & McNerney, 2011; Rothen & Meier, 2010a; Yaro & Ward, 2007), *calendar-form synaesthetes*, who experience dates as located in peripersonal space, have advantages for remembering events and dates (Simner, Mayo, & Spiller, 2009), and several varieties of synaesthesia are associated with enhanced perceptual discrimination (Banissy, Walsh, & Ward, 2009; Saenz & Koch, 2008). However it remains an open question

5. This chapter is slightly adapted from a previously published paper (Watson, Blair, Kozik, Akins, & Enns, 2012). I was the primary person involved in determining the research question and experimental method, though all my collaborators made many useful suggestions and changes. All analyses were performed by myself, and I was the primary author of the paper.

whether synaesthesia can be exploited for more sophisticated and abstract forms of learning (Brang & Ramachandran, 2011). Here we answer this question in the affirmative for rule-based category learning.

We investigated whether grapheme-color synaesthetes are able to use synaesthetic colours on a difficult category learning task. We show that synaesthetes viewing black letters use their internally-generated colours during this task in much the same way as non-synaesthetes viewing genuinely coloured stimuli. Thus synaesthesia can be a tool used in learning novel abstractions.

Participants learned to classify stimuli according to a rule-based category structure. Such learning is hypothesized to involve an explicit reasoning process in which hypotheses are maintained in working memory, individual stimuli are attended to and categorized according to the currently active hypothesis, and this hypothesis is either strengthened or modified on the basis of subsequent feedback (Ashby & Maddox, 2005). The particular 4-category structure we created was structurally similar to one used by Maddox, Filoteo, Hejl, and Ing (2004), in that the category rules conjoin two distinct pieces of information. Such conjunctive rules are frequently taught in primary school, for example when learning English phonetics (e.g. a vowel followed by a consonant has a short pronunciation, unless the consonant is immediately followed by the letter 'e', in which case the first vowel has a long pronunciation), in mathematics (e.g. a number is prime if it can be divided by 1 and not by any other number), or in the sciences (e.g. a mammal is an animal with warm blood that gives birth to live young).

Stimuli were pairs of graphemes (see Figure 6.1a) whose category membership could be determined by simple rules involving the order and associated colours of graphemes, e.g. ''Members of category 1 contain a green followed by a pink grapheme''. As synaesthetes' colours are idiosyncratic, a different stimulus set was generated for each synaesthetic participant. Participants who discovered the colour rules were expected to be more accurate on the initial Category Learning task than those who did not. Other participants would have to resort to more complex rules based on all possible combinations of the eight graphemes in the stimulus set, to use explicit memorization of 16 stimulus-category pairs, or to resort to more idiosyncratic strategies – e.g. treating letter pairs as acronyms for words or phrases with personal meaning.

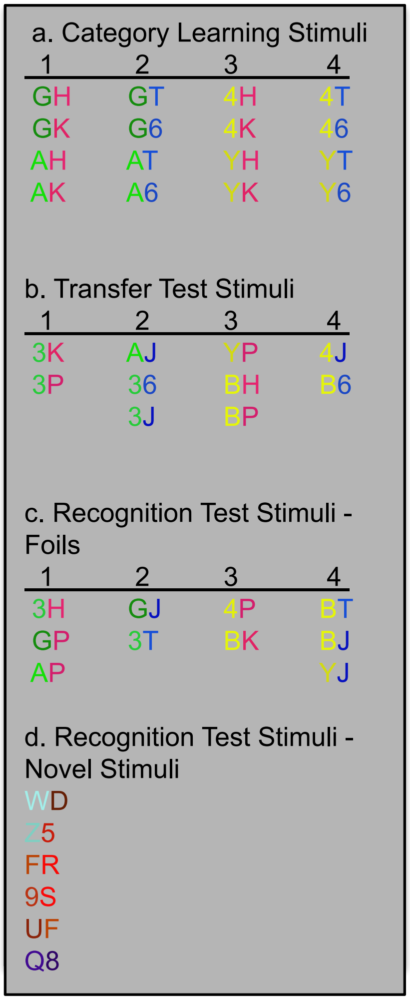

Figure 6.1 One of the stimulus sets, based on the color assignments of one of the synaesthetes, for the various measures in the experiment. Synaesthetes and members of the Control-Achromatic group would have been presented with these stimuli in black. (a) Stimuli used during the Category learning task and Recognition Test, arranged so that the color rules are obvious. (b) The 10 stimuli used during the Transfer Test. (c) The 10 Foil Stimuli used during the Recognition Test. (d) The six completely novel stimuli used during the Recognition Test.

This initial task was followed by a Transfer Test and Recognition Test, both designed to verify if participants were using colour rules. Immediately following the category learning task, participants completed 10 Transfer Test trials, in which novel stimuli that followed the same colour rules were presented (Figure 6.1b). Participants who had used colour rules previously ought to be able to apply them to these novel stimuli, whereas those who used other strategies should be at chance. This Transfer Test was followed by a Recognition Test, on which the opposite pattern of results was expected. Here participants were presented with grapheme pairs and asked if they had been presented previously in the experiment. These stimuli included all 16 stimuli from the Category Learning task (Figure 6.1a), 10 novel Foil Stimuli that also followed the colour rules (Figure 6.1c), and six additional stimuli with no colours or identities in common with any others used during the experiment (Figure 6.1d). Subjects who had used colour rules would be expected to confuse the 10 Foil Stimuli with those previously presented in the category learning task and Transfer Test. However those using alternative rules would be expected to correctly reject more of them.

Three groups participated in the study. A Synaesthete group viewing achromatic stimuli was compared with non-synaesthetes viewing either the same achromatic stimuli (Control-Achromatic) or stimuli that were coloured according to synaesthetic colour assignments (Control-colour). Thus if synaesthetic colours can be used in rule-based categorization tasks, we expect the Synaesthete group to perform better than the Control-Achromatic group on the category learning task and the Transfer Test, but worse on the Recognition Test. Comparing the Synaesthete and Control-colour groups allows us to infer further similarities and differences between synaesthetic and normal colour perception.

6.2 Experiment 1

6.2.1 Participants

Ten grapheme-colour synaesthetes participated in the study and were rewarded with $10 (CAN). All synaesthetes' grapheme-colour associations were verified as consistent by the online Synaesthesia Battery (Eagleman, Kagan, Nelson, Sagaram, & Sarma, 2007), with a mean consistency score of .70, and a mean accuracy score of 89% on the Speed-Congruency Test. Eighty- six non-synaesthetes were recruited from undergraduate psychology classes at the University of British Columbia. Six of these participants were removed from all analyses for performing at chance, leaving 80 non-synaesthetic participants. Eight of these participants were randomly assigned to each synaesthete's stimulus set, 4 to the Achromatic and 4 to the colour condition.

6.2.2 Displays and responses

Stimuli consisted of grapheme pairs presented in Arial font, each grapheme occupying approximately 2.5 cm^2 (3.6° of visual angle at a distance of 40 cm). Graphemes were either all black (Achromatic condition) or coloured as the synaesthete reported them (colour condition). Each Category Learning and Transfer Test trial began with a fixation cross (approximately 1.5 cm^2, or 2.1° v.a.) at the center of the screen for 400–800 ms, followed by a stimulus at the center of the screen and four response boxes near the bottom, labeled with the digits 1–4. Participants selected one box with a mouse click, and were given feedback in the form of the incorrect response boxes disappearing. Participants responded to the feedback by clicking on the correct box, and the next trial began. Recognition Test trials had identical displays, save that there were only two re- sponse boxes, labeled ''Yes'' and ''No'', and no performance feedback was given.

6.2.3 Category structure

A category structure of 16 letter pair stimuli was created for each of the 10 synaesthetes. These were generated from eight graphemes, which were associated with four distinct colours (see Figure 6.1a), organized such that each colour appeared only in the left or right position. Stimuli were assigned to one of four categories on the basis of simple conjunctive colour rules, as illustrated in Figure 6.1a that can be easily applied in a 2-stage hierarchy. For example, one could begin a trial by looking at the left- hand letter, and narrowing down the possible responses to categories 1 and 2 if the letter is blue or 3 and 4 if it is red. Then the colour of the right hand letter could be used to determine which of the two remaining options is correct, since the possible responses are 1 and 3 if this letter is orange or 2 and 4 if it is green. Of course these colour rules would be unavailable to non- synaesthetes viewing achromatic letters, and we expected their performance to suffer accordingly.

6.2.4 Foil stimulus sets

In addition to the stimulus sets shown to participants during Category Learning, we constructed stimulus sets following the same colour rules, but including novel graphemes, for use in the subsequent Transfer and Recognition Tests, illustrated in Figure 6.1b and c. Within each set, four additional graphemes were used, each associated with one of the four colours from the learning phase. Combined with the original graphemes, this allowed for the construction of 20 more grapheme pairs (five new stimuli in each category) that followed the same colour mapping as in the learning task. Ten of these stimuli were randomly selected to appear in the Transfer Test, and the other 10 appeared during the Recognition Test. Again, participants in the Achromatic condition saw the same letter pairs, but coloured black.

6.2.5 Procedure

Category Learning consisted of 256 trials in total, divided into eight blocks of 32 trials. Each block contained each of the 16 stimuli presented twice in random order. On each trial participants indicated which category a stimulus belonged to and were given feedback as described above. Immediately following these eight blocks, participants completed 10 Transfer Test trials that were identical in format, except that the stimuli were drawn from the Foil Stimuli. Other than the sudden appearance of novel stimuli, participants were given no indication that anything was different on these trials.

Participants then completed 32 Recognition Test trials, where they were asked to indicate if they had seen a particular stimulus previously during the experiment. They indicated their response by clicking on one of two boxes, labeled ‘‘Yes’’ and ‘‘No’’, and were not given any feedback. The stimuli presented in this phase consisted of all 16 original stimuli, the 10 Foil Stimuli that had not been used in the Transfer Test, and six new grapheme pairs unrelated to any of the other stimuli in the experiment. Thus, half of the stimuli in the Recognition Test had been seen previously and half had not. We were particularly interested in participants' responses for the 10 Foil Stimuli, as someone paying attention to colour might be expected to make False Recognition errors on these trials.

Finally, participants were asked to write down any strategies they used during the Category Learning phase of the experiment.

6.2.6 Behavioral results

For each participant, we computed four scores. Categorization Accuracy was the mean accuracy over each block of the category learning task, and response times (RTs) were also recorded during these blocks. Transfer Accuracy was the mean accuracy over the 10 Transfer Test trials. False Recognition was the inverse of the mean accuracy over the 10 Recognition Test trials that used the Foil Stimuli. (Recognition Test accuracy for the other stimuli was over 95% for all groups, and so was not analyzed further.)

The results were qualitatively very simple. First, accuracy on the Category Learning task was higher for those with access to colour information, whether these colours were synaesthetic or real (see Figure 6.2a), although synaesthetes learned somewhat more slowly than controls viewing real colours. Second, participants looking at achromatic letters were slower to make decisions, whether they were synaesthetes or controls, and the synaesthetes were generally slowest of all. Third, access to colour information also improved participants' ability to generalize to novel stimuli on the Transfer Test, although real colours provided more of an advantage than synaesthetic colours (see Figure 6.2b). Finally, participants with access to colour information were prone to False Recognition of the Foil Stimuli during the Recognition Test, but those without colour were able to correctly re- ject most of these stimuli (see Figure 6.2c).

These qualitative descriptions are supported by analyses of variances (ANOVAs) and post hoc group comparisons. To begin with, Categorization Accuracy and RT were the dependent measures in two-way ANOVAs using Group as a between-subjects factor with three levels (Synaesthete, Control-Achromatic, and Control-colour), and Epoch (1–4, each composed of two experimental blocks) as a within-subjects factor. In both cases, there were significant main effects of Group (Categorization Accuracy: $F_{2,87}$ = 12.1, η^2 =

.22, MSE = .11, p < .001; RT: $F_{2,87}$ = 11.5, η^2 = .21, MSE = 2.1, p < .001) and Epoch ($F_{3,261}$ = 146.5, 31.4, η^2 = .61, .26, MSE = .01, .42, respectively; both ps < .001), as well as Block by Epoch interactions ($F_{6,261}$ = 2.5, 2.3, $\eta 2$ = .02, .04, MSE = .01, .42, respectively; both ps < .05). These were followed by tests of the simple main effect of Group at each of the four levels of Epoch, which all indicated group differences ($F_{2,261}$ between 5.0 and 19.3, $\eta 2$ between .04 and .12, MSE for accuracy between .06 and .14, for RT between 2.1 and 8.1, all ps < .01; except for RT on epoch 4, where $F_{2,261}$ = 3.2, $\eta 2$ = .02, MSE = 1.3, p = .04). The Tukey–Kramer method was used to determine which groups were significantly different from each other on each of the four epochs, and these results are described below.

In the case of Categorization Accuracy, as shown in Figure 6.2a, the interaction stems from the Synaesthete group improving at a faster rate (a rise of over 40% from epochs 1 to 4) than either control group (both of whom improve by approximately 30%). The Control-colour group outperforms the Control-Achromatic group by 15–20% throughout the experiment, while the Synaesthete group begins by performing similarly to the Control-Achromatic group on the first epoch, but on Epochs 2–4 is significantly more accurate than the Control-Achromatic group, and not distinguishable from the Control-colour group.

In the case of RT, the interaction also stems from the Synaesthete group improving at a faster rate (an overall gain of 1.7 s from an initial RT of 4.9 s in Epoch 1) than either control group (the Control-colour and Control-Achromatic groups improve by 0.6 s from 2.1 s and by 0.8 s from 2.8 s, respectively). Despite this greater improvement, the Synaesthete group was slower than both control groups on all epochs save for epoch 4, where it was not distinguishable from the Control-Achromatic group. The Control-colour group was faster than both other groups on all epochs save for epoch 3, where it was not distinguishable from the Control-Achromatic group.

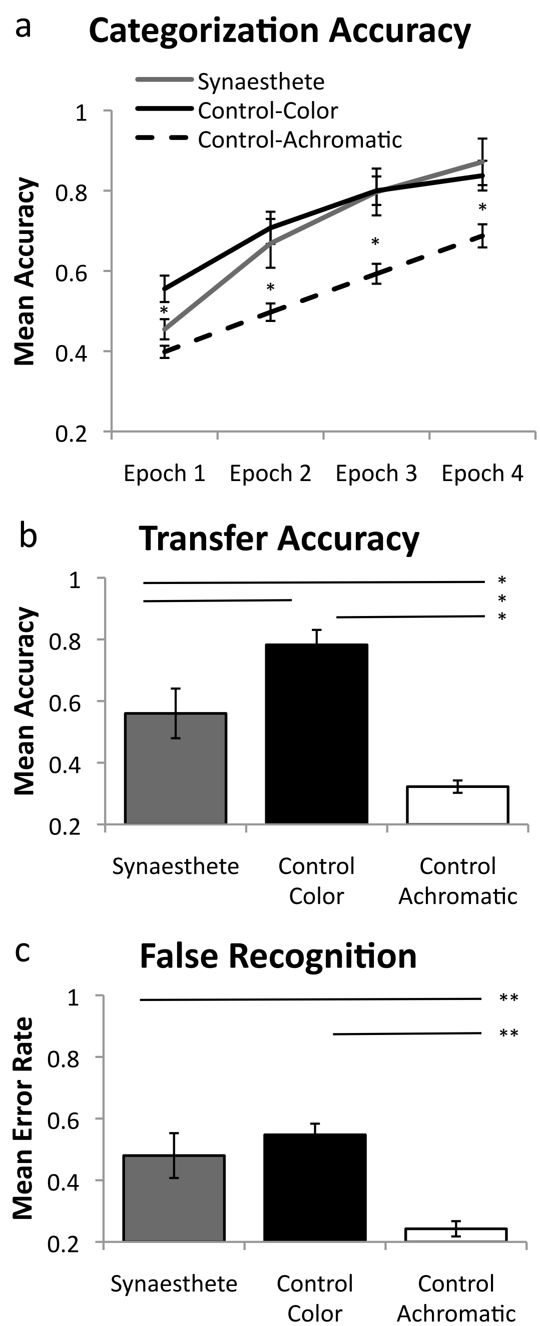

Figure 6.2 Performance of participants in Experiment 1. (a) Accuracy over the course of the Category Learning task (each epoch = 64 trials), (b) accuracy over the 10 Transfer Test trials in which participants categorized foils that follow the same color mapping rules, and (c) error rate over the 10 Recognition Test trials involving novel Foil Stimuli. Error bars indicate plus/minus one standard error of the mean. Asterisks indicate significant group differences. *: $p < .05$, **: $p < .01$.

The remaining two measures (Transfer Accuracy and False Recognition) were dependent variables in one-way ANOVAs using Group as a between-subjects factor. Levene's test showed a violation of the homogeneity of variance assumption for Transfer Accuracy ($F_{2,87}$ = 4.5, p < .05) so Welch's statistic was used. Both ANOVAs were significant (Transfer Accuracy: $F_{2,22.2}$ = 39.4, η^2 = .46, *MSE* = .06, p < .001; False Recognition: $F_{2,87}$ = 24.2, η^2 = .36, *MSE* = .04, p < .001), indicating group differences on each of the measures. Following up with Tukey–Kramer revealed that the Control-colour group had the highest Transfer Accuracy (78%), then the Synaesthete group (56%), followed by the Control-Achromatic group (32%), and all 3 between-group comparisons were significant (all ps < .05). Finally, the Synaesthete and Control-colour groups both performed poorly on the Recognition Test trials using Foil Stimuli (False Recognition of 48% and 55%, respectively) whereas the Control- Achromatic group made far fewer errors (24%) than either (both ps < .01).

6.2.7 Self-report data

The data from participants' reports of their own strategies also support the notion that group performance differences stemmed from the availability of colour-based rules. Reviewing these reports revealed a number of common strategies, including the use of acronyms (mentioned by 14% of participants), memorization (60%), various forms of mathematical reasoning (11%), the use of colour information (49%), and explicit descriptions of the formal category structure (22%). Fisher's Exact Test was used to see if the proportions of participants reporting each strategy differed between groups. This was the case for memorization (p = .002), the use of colour (p < .001) and describing the structure (p < .001), but not for acronyms (p > .9) or math (p > .25). The Control-Achromatic group was the source of all three group differences, as removing this group resulted in no significant differences between the Synaesthete and Control-colour groups (all *p*s > .5). Specifically, the Control Achromatic group was more likely to report memorization (80%) than the Synaesthete (50%) or Control-colour (43%) groups, and less likely to report the use of colour (0% vs. 80% and 90%, respectively) or to describe the category structure (3% vs. 40% and 38%, respectively).

Finally, we verified whether these strategies were connected to performance using five ANOVAs, each using Group and one of the strategies described above as between-subjects factors, and accuracy on the final category learning epoch as the dependent measure. There were main effects of describing the category structure ($F = 7.1$, $\eta^2 = .08$, $MSE = .04$, $p < .01$) and mathematical reasoning ($F = 6.8$, $\eta^2 = .05$, $MSE = .04$, $p < .01$), and an interaction between the use of colour information and group ($F = 5.2$, $\eta^2 = .05$, $MSE = .04$, $p < .05$). No other main effects of strategy or interactions were significant (all $ps > .05$). The two main effects of strategy were due to participants who described the category structure performing better than those who did not (mean accuracy on Epoch 4 = 94% vs. 73%) and those who used mathematical reasoning performing worse than those who did not (60% vs. 80%). The group by colour interaction was followed with tests of the simple main effect of colour, which was marginally significant for the Control-colour group ($F = 3.8$, $\eta^2 = .04$, $MSE = .15$, $p = .06$), but not for the Synaesthete group ($p > .25$). Among Control-colour participants, Tukey–Kramer revealed that participants who used colour had a higher accuracy than those who did not (87% vs. 59%, respectively). As only two synaesthetes did not report using colour, the lack of a main effect is likely uninformative for this group.

6.3 Experiment 2

The overall pattern of results from Experiment 1 is consistent with the claim that synaesthetes can exploit their colour experiences during category learning, in much the same manner as non-synaesthetes viewing real colours. Indeed, the only accuracy differences between the Synaesthete and Control-colour groups were that the synaesthetes were slightly slower to learn the category structure, and somewhat less accurate on the Transfer Task. But it was still possible that the superior performance of the Synaesthete and Control-colour groups was not due to their using colour rules per se. For instance, it is possible that colour by itself makes the categorization task easier to learn, irrespective of any category rules: perhaps it is simply easier to memorize letter pairs when they are coloured, and hence easier to apply mnemonic strategies to learn the category structure. To see if this was the case, we ran a second experiment using new subjects, identical to Experiment 1 save that colours were no longer diagnostic of category membership. If the results from Experiment 1 were indeed due to participants in the Synaesthete and Control-colour groups making category decisions on the basis of colour rules, then they should not be able to do this in Experiment 2, and we should not find the group differences we found in Experiment 1.

6.3.1 Methods

The experimental procedure was identical to Experiment 1, save that similarly-coloured letters were not grouped in the same categories, so colours were no longer diagnostic of category membership. Eight synaesthetes participated in the experiment, whose grapheme-colour associations were verified as consistent by the Synethesia Battery (mean consistency score: .66, mean Speed-Congruency accuracy: 87%), along with 68 non-synaesthetic controls, four of whom were eliminated from the analysis as random responders, leaving 64 non-synaesthetes who were randomly assigned to a particular synaesthete's stimulus set, and to a colour or Achromatic condition, as in Experiment 1. None of the participants were in Experiment 1.

As in Experiment 1, stimuli were composed of pairs of graphemes, made from eight graphemes with four distinct colours. These were organized into four categories of four stimuli each. However graphemes with each of the four colours appeared in at least three of the four categories, and at least once on the left and once on the right-hand side of different stimuli. Thus colour was entirely useless for categorization.

6.3.2 Results

In brief, the three groups perform similarly to each other on all tasks, with only one exception. Furthermore, all three groups also perform very similarly to the Control-Achromatic group from Experiment 1. Thus it is clear that the group differences we found in Experiment 1 are not due to colour per se, but to its use in a system of rules.

To support this conclusion, we performed the same analyses on the same variables as in Experiment 1. No group differences or interactions were found (all ps > .3) save for False Recognition ($F = 5.4$, $\eta^2 = .14$, $MSE = .05$, $p < .01$). Tukey's HSD revealed that this group difference was due to the Control-colour group performing significantly worse ($p < .01$) than the Control-Achromatic group (False Recognition of 57% and 28%, respectively). Furthermore, with the exception of the Control-colour group's False Recognition, all groups' mean performance on all measures was within the 95% confidence interval of the performance of the Control-Achromatic group on Experiment 1.

6.4 Discussion

These results demonstrate that synaesthetes can learn rule-based categories using internally-generated synaesthetic colours. Moreover, they do this similarly to non-synaesthetic individuals using physical colours. Both synaesthetes and non-synaesthetic participants viewing coloured stimuli learned to categorize more successfully than non-synaesthetes viewing achromatic stimuli, were able to generalize to novel stimuli on the transfer task, and were unable to correctly reject Foil Stimuli in a Recognition Test, indicating that their memory for individual grapheme identities was impaired. Furthermore, these participants were also likely to give explicit reports indicating that they used the colour information and understood the category structure, unlike non-synaesthetes viewing achromatic stimuli, and giving these reports was correlated with higher accuracy. Taken together, these findings demonstrate that synaesthetes can exploit their grapheme colours to learn a rule-based category structure similar to those taught in a variety of domains.

More detailed analyses showed some performance differences between synaesthetes and non-synaesthetes viewing physically coloured stimuli. First, synaesthetes learned more slowly. Though their performance for most of the experiment was comparable to non-synaesthetes viewing real colours, their accuracy was lower at the start of the experiment. We suggest that this is because experiences of synaesthetic colours may be somewhat less vivid than experiences of real colours, which might delay rule acquisition.

Second, synaesthetes were not as successful in transferring their learning to novel stimuli. This might also be explained by less vivid synaesthetic experiences. Alternatively, a comment made by a synaesthetic participant may shed light on this result. He indicated that when viewing the stimuli, he did not experience two different colours, but saw a single colour for the pair as a whole, typically the colour of the grapheme that seemed more ‘‘dominant’’ than the other. Indeed, many grapheme-colour synaesthetes experience single colours for words, often determined by the colour of an individual letter (Simner, Glover, & Mowat, 2006; Ward, Simner, & Auyeung, 2005). This may account for the lower accuracy of synaesthetes on the transfer task, although it does not mitigate the critical finding that their accuracy was almost twice that of non-synaesthetes viewing achromatic grapheme pairs.

Third, synaesthetes were slower to respond than participants viewing real colours. There are at least two ways of accounting for this result. First several researchers argue that synaesthetic colours cannot be induced without the conscious recognition of the grapheme (e.g. Laeng, 2009). This would imply that the Synaesthete group ought to respond at least as slowly as the Control-Achromatic group, which is what we find. An alternative is that the process of establishing which letter in a pair is dominant, as described in the previous paragraph, may take some time to resolve itself. The present data does not provide enough evidence to decide whether one or both of these is the true source of the reaction time differences.

How well might these results generalize to other tasks? Stimuli in Experiment 1 were specifically tailored to each synaesthete such that their personal colour associations would be maximally informative for distinguishing between the four categories. It seems remarkably unlikely that this could happen by chance. Thus one would be justified in asking whether our results have any meaning outside the laboratory. Are the apparently arbitrary associations synaesthetes make between graphemes and colours actually any use in learning or using the rules of, for example, spelling, mathematics or phonetics?

Our data do not directly address this question, but there is reason to think that synaesthetic colours could provide a significant benefit to such rule use. Many of the explicit rules we learn in everyday life – including all the examples given in Section 1 – are single rules that do not require combining with other rules in a hierarchical fashion, as the colour rules in this study do. Any rule that involves a specific combination of letters – e.g. ''*I* before *E* except after *C*'' – is one that involves a specific combination of colours for a grapheme-colour synaesthete. Provided the synaesthete's colours for these letters are distinguishable from each other, this could provide a cue to aid in learning and applying the rule. The same is true for any numerical rule in math – e.g. any number ending in 5 is divisible by 5. The present study shows that with less than 30 min training, synaesthetes can flexibly employ their colours to learn and use a complex and abstract set of intertwined rules. We see no reason why they could not do the same for simpler rules in the classroom or in the rest of daily life.

Of course establishing that this is possible is one thing, verifying that it occurs under natural conditions is another matter. There are no published studies that test this hypothesis. There are anecdotal reports, and several savants attribute their astounding memory and mathematical skills to synaesthesia (Bor, Billington, & Baron-Cohen, 2007; Luria, 1968), but it remains unconfirmed whether the average synaesthete employs their colours in this manner. It appears that synaesthetic photisms influence mathematical processing (Ghirardelli, Mills, Zilioli, Bailey, & Kretschmar, 2010), but the nature of this influence is far from clear. Determining if colours are actually being used to represent rules in mathematics, spelling, or in other domains is a crucial next step.

Finally, the rule-based categorization task used here is generally considered to involve explicitly conscious processes that operate in a fundamentally different manner from the processes used in statistical or implicit learning (Reber, 1993). Further experiments could more directly test whether the sub-personal mechanisms that underlie implicit learning can also exploit synaesthetic colour information. If this is the case, the potential utility of synaesthesia for learning is even wider, given the ubiquity of implicit learning throughout life.

Previous work has demonstrated that the colours synaesthetes associate with letters are influenced by a number of learned properties of these letters (Beeli, Esslen, & Jäncke, 2007; Cohen Kadosh, Henik, & Walsh, 2007; Day, 2005; Rich, Bradshaw, & Mattingley, 2005; Simner et al., 2005; Simner & Ward, 2008; Watson, Akins, & Enns, 2012). Here, we demonstrate the reverse: synaesthetic colours can influence learning about letters. Further exploration of the interactions between synaesthesia and learning is likely to be the source of new understanding about the nature of this fascinating phenomenon.

7 Conclusion

7.1 What do we know?

7.1.1 Demonstrating 3 main points

The work described in this thesis was inspired by the developmental learning hypothesis of synaesthesia, according to which synaesthesia develops, at least in part, as a strategic method of solving various learning challenges in childhood. The thesis set out to defend three points which underpin this hypothesis. These are:

1. Synaesthesia typically develops as part of a difficult learning process in which the synaesthete learns a category structure whose members become the triggers of synaesthetic experiences.
2. Synaesthetic experiences are shaped by this learning process, such that they encode a wide range of information about the learned domain.
3. Synaesthesia is exploited on a variety of memory, learning, and creative tasks, leading to “synaesthetic styles” of performance on these tasks.

In the Introduction, I presented what we already know about these three points, and in the research chapters I expanded on this by presenting new evidence supporting them.

For point 1, the literature already shows that the majority of synaesthetic inducers are learned categories, and that the development of grapheme-colour synaesthesia takes place over a period of several years that overlap with the development of literacy. Chapter 2 expanded on this, by showing that a particular learning challenge—the acquisition of a second language in grade school—is associated with a higher likelihood of becoming synaesthetic.

Point 2 is supported by a number of previous findings demonstrating that synaesthetic concurrents are shaped by a wide range of learned information about the inducer domain. In Chapters 3-5, I expanded on this, showing that various different kinds of learned information independently influence the colours of grapheme-colour synaesthetes. In Czech, this includes quite high-level categorical knowledge about the role different letters and diacriticals play within the language. This knowledge is not available until first grade at the earliest, showing that synaesthetic development is mediated by high-level learning over the course of several years.

Finally, point 3 has been largely accepted by researchers, but generally supported by anecdotal evidence. Chapter 6 provides the first clear proof that synaesthetic associations can be useful on a difficult learning task. Synaesthetes in this category learning study spontaneously began using their colours as guides to the category structure in a very similar manner to non-synaesthetes looking at real colours, and both groups significantly out-performed non-synaesthetes who did not have access to colour information. If synaesthetes can learn to do this in 20-40 minutes seated in front of a computer, who knows what they can do over the course of decades as they navigate the various tasks associated with literacy?

7.1.2 Other findings

In the course of gathering evidence for the three main points guiding this thesis, my collaborators and I uncovered a number of other fascinating findings, which may or may not directly relate to the developmental learning hypothesis. Here I outline some of the most interesting of these findings.

Chapter 2 demonstrated that women are far more likely to cooperate with experimenter requests in synaesthesia studies, demonstrating that previously-reported female biases in synaesthesia are almost certainly due to cooperation with the experimental protocol, rather than any actual differences in prevalence between the sexes. It is probably now safe to assume that there is at most a very small influence of sex on synaesthetic prevalence, and very likely no influence at all. If there is no influence, this suggests that the female compliance bias extends all the way to answering direct questions about synaesthesia: men may be less likely to self-report as synaesthetic even if they are given a description of synaesthesia and asked if this applies to them. If this is the case, then all previous reports of sex differences in synaesthesia should be treated very cautiously, in particular the finding of a strong bias for transmission down the maternal line. This casts doubt upon one of the main reasons for believing in a simple genetic account of synaesthetic development, which in turn makes something like the developmental learning hypothesis far more plausible.

Another important result from Chapter 2 is the finding of a large proportion of individuals who pass the consistency tests for various kinds of synaesthesia, but initially denied having the forms of synaesthesia in question. This confirms that the self-report and consistency criteria for synaesthesia are fully independent of each other, and shows that we cannot simply screen synaesthetes based on self-report. In the Introduction I reviewed the current consensus that we do not have a good theoretical definition of synaesthesia, and these data suggest that there are serious issues with our operationalizations as well.

Chapter 2's results are also consistent with previous reports of synaesthetic "clustering". We find that the effect of second language acquisition is not confined to grapheme-colour synaesthesia, as originally hypothesized, but rather to all varieties of synaesthesia we test for. This, it is argued, may be the result of generalizing a particular learning strategy initially developed during literacy acquisition to other, conceptually similar, problems.

One fascinating finding from Chapters 3-5 is that almost all the effects we uncover are confined to either luminance or hue, but not both. In Chapter 3 it was argued that this split is useful. Categorical information can be most usefully mapped onto hue space, partly because this is more natural for humans, but also because this would not interfere with the luminance-based processes we use to actually identify letter categories. This suggestion is complicated somewhat by the Czech findings from Chapter 5. While the same hue/luminance splits were found for several mappings, the strongest categorical mappings (the base-diacritical pairs) were found in both luminance and hue, as was the influence of phonology.

The Czech data used in Chapter 5 also enabled a strong test of the learning order hypothesis, which stated that the ordinality-hue effect, in which letters earlier in the alphabet are further apart in hue, was due to the earlier letters having been learnt first. Since Czech letters are not learned in alphabetical order, this hypothesis could be directly tested, and it was conclusively rejected. Instead, it appears that there is something special about ordinality itself, which may provide support for Novich *et al.*'s notion of coloured sequence synaesthesias.

The data from Chapter 5 also allowed a tentative ranking of some of the influences on synaesthetic colour, suggesting that letter shape and alphabetical order (and possibly an abstract notion of letter identity) are stronger determinants of colour than phonology. This might be taken as evidence that grapheme-colour synaesthesia does not have its earliest roots in a phoneme-colour synaesthesia, but rather results specifically from the learning problems associated with the development of literacy.

Finally, in Chapter 6 we saw that though the synaesthetes in our study exploited their colours in a similar manner to non-synaesthetes viewing real colours, their response times were far slower. This may be an indication that the induction of synaesthetic experiences is a slow process, which might be taken as support for the re-entrant processing or disinhibited feedback theories of the neurophysiology of synaesthesia. As always, more investigation is needed.

7.2 What don't we know?

I take it that the three points that were the organizing principles of this thesis have been firmly established. This does not, it should be noted, prove the developmental learning hypothesis *per se*. In order to do this, one would need to connect the points in the causal chain described in the Introduction, which is far beyond the scope of this thesis. That is, one would have to show that synaesthesia develops in the course of learning (point 1) *because* it can be exploited to achieve learning success (point 3), and that the particular aspects of the inducer domain which are encoded within synaesthetic concurrents (point 2) are encoded because they are useful in this manner. I hope that this causal chain seems at least somewhat plausible now. For example, it seems intuitive that encoding categorical information about vowels and consonants might be particularly useful in learning to read. However plausibility is a long way from proof, and a great deal of innovative research would need to take place before this causal chain could be established. What might this research look like?

One step would be to verify if children with synaesthetic colours are exploiting them in some way. Consider the case of the categorical distinction between vowels and consonants. Would a child whose synaesthetic colours make this distinction clear be impaired on a vowel/consonant discrimination task if the letters were presented in incongruent colours? This would at least provide preliminary evidence that their method of making this discrimination in some way exploited their colours.

A great deal of reading performance involves the recognition of dozens of common bigrams and trigrams. Colour-based “markers” for letters might be a particularly good method of recognizing such common groups of letters, allowing for improved acquisition and retention. Another possible study, then, would be to see, for individual synaesthetes, which of the common bigrams and trigrams are particularly salient, given their personal letter colours, and then to see if they have a particular advantage at recognizing, discriminating or otherwise interacting with these groups of letters above others.

An important longitudinal question concerns what happens to those children who at some point have synaesthetic tendencies, but who lose these as they get older. Might these be individuals for whom letter colours were particularly uninteresting, or for whom the development of literacy was accomplished in a manner that did not exploit their synaesthesia?

Finally, the research presented here has been strongly slanted towards grapheme-colour synaesthesia, and letter-colour synaesthesia in particular. However we have seen strong evidence that this emphasis does not reflect the tendencies among synaesthetes, who generally have multiple forms of synaesthesia. Understanding how synaesthetic colours (or other concurrents) might help with learning these other inducer domains is an important and difficult task.

7.3 Why study synaesthesia?

Supposing the developmental learning hypothesis is true, one can imagine a great number of ways in which the kinds of studies described here might have practical benefits. For instance, if colour-coding information is useful for synaesthetes, might it be equally useful for non-synaesthetes? Would it help with second-language learning among older children or adults? Anything that could reduce the drudgery and increase the success rate of second-language learning would be of obvious utility to many. Other similar possibilities are fairly easy to think of.

Aside from its potential practical utility, it is often suggested that synaesthesia is of great theoretical importance. Researchers often describe the importance of synaesthesia research in terms of its utility for understanding the mysteries of consciousness, multimodal associations, the heterogeneity of perception across individuals, or other profound aspects of the human mind. One question which might have real theoretical importance, for example, concerns the great preponderance of spatial and coloured concurrents in synaesthesia. Though there are synaesthetes who taste words, or hear sounds when they see particular stimuli, they are a tiny minority compared to those who experience members of categories as coloured or as located in personal space. Why are space and colour so important? Understanding this might be of real importance to understanding the structure of our own minds, which is arguably the purpose of psychology.

However, while I agree that such theoretical and practical considerations are important, they are not the ones that motivated me to begin studying synaesthesia, and I suspect this is true of most researchers. I study synaesthesia for its own sake. It is seriously weird, and endlessly fascinating. I argue that this is a perfectly good reason to devote time and energy to a research topic. Basic research cannot be guided purely by short-term utilitarian concerns, *especially if a society's long-term interests are utilitarian*. By definition, unexpected results that genuinely shift paradigms and provide the greatest long-term benefits to society cannot come from "safe" areas of research, because all that we mean by "safe" is that an area of research is expected to produce useful results.

Frankly, while I agree that synaesthesia research can shed light on consciousness or other important issues, it is not at all clear to me that it is more likely to do this than research into virtually any other conscious phenomenon. And while colour-coding learning materials might improve retention, I suspect that raising teacher's salaries would have a far stronger effect. However science needs to keep poking into the weird and fascinating, because every now and then when we turn over one of these pebbles, we find it is actually a diamond.

Bibliography

Albright, A. (2006). SimilarityCalculator.pl. Retrieved from http://web.mit.edu/albright/www/software/SimilarityCalculator.zip

Aleman, A., Rutten, G. J. M., Sitskoorn, M. M., Dautzenberg, G., & Ramsey, N. F. (2001). Activation of striate Cortex in the absence of visual stimulation: an fMRI study of synesthesia. *Neuroreport*, *12*(13), 2827-2830.

Amin, M., Olu-Lafe, O., Claessen, L. E., Sobczak-Edmans, M., Ward, J., Williams, A. L., & Sagiv, N. (2011). Understanding grapheme personification: A social synaesthesia? *Journal of Neuropsychology*, *5*, 255-282.

Anonymous. (2011). Kanye West behind the scenes interview. *YouTube*. Retrieved March 15, 2013, from http://www.youtube.com/watch?v=UhXKhSer2ic

Anonymous. (n.d.). List of superhero characters with synaesthesia powers. *Comic Vine*. Retrieved March 12, 2013, from http://www.comicvine.com/synaesthesia/41-129/

Asano, M., & Yokosawa, K. (2011). Synesthetic colors are elicited by sound quality in Japanese synesthetes. *Consciousness and Cognition*, *20*(4), 1816-1823.

Asano, M., & Yokosawa, K. (2012). Synesthetic colors for Japanese late acquired graphemes. *Consciousness and Cognition*, *21*(2), 983-993. doi:10.1016/j.concog.2012.02.005

Ashby, F. G., & Maddox, W. T. (2005). Human category learning. *Annual Review of Psychology*, *56*, 149-178.

Asher, J. E., . . . Monaco, A. P. (2009). A Whole-Genome Scan and Fine-Mapping Linkage Study of Auditory-Visual Synesthesia Reveals Evidence of Linkage to Chromosomes 2q24, 5q33, 6p12, and 12p12. *American Journal of Human Genetics*, *84*(2), 279-285.

Bailey, M. E. S., & Johnson, K. J. (1997). Synaesthesia: Is a genetic analysis feasible? In S. Baron-Cohen & J. E. Harrison. Oxford: Basil Blackwell.

Banissy, M. J., Stewart, L., Muggleton, N. G., Griffiths, T. D., Walsh, V. Y., Ward, J., & Kanai, R. (2012). Grapheme-color and tone-color synesthesia is associated with structural brain changes in visual regions implicated in color, form, and motion. *Cognitive Neuroscience*, *3*(1), 29-35.

Banissy, M. J., Walsh, V., & Ward, J. (2009). Enhanced sensory perception in synaesthesia. *Experimental Brain Research*, *196*(4), 565-571.

Barnett, K. J., Finucane, C., Asher, J. E., Bargary, G., Corvin, A. P., Newell, F. N., & Mitchell, K. J. (2008a). Familial patterns and the origins of individual differences in synaesthesia. *Cognition*, *106*(2), 871-893.

Barnett, K. J., Foxe, J. J., Molholm, S., Kelly, S. P., Shalgi, S., Mitchell, K. J., & Newell, F. N. (2008b). Differences in early sensory-perceptual processing in synesthesia: A visual evoked potential study. *Neuroimage*. doi:10.1016/j.Neuroimage.2008.07.028

Baron-Cohen, S. (1996). Is there a normal phase of synaesthesia in development? *Psyche: An Interdisciplinary Journal of Research on Consciousness*, *2*(2).

Baron-Cohen, S., Burt, L., Smith-Laittan, F., Harrison, J., & Bolton, P. (1996). Synaesthesia: Prevalence and familiality. *Perception*, *25*(9), 1073-1079.

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

Baron-Cohen, S., & Harrison, J. E. (Eds.). (1997). *Synaesthesia: Classic and Contemporary readings*. Oxford: Basil Blackwell.

Baron-Cohen, S., Wyke, M. A., & Binnie, C. (1987). Hearing words and seeing colours: an experimental investigation of a case of synaesthesia. *Perception*, *16*(6), 761-767.

Education, B. C. M. o. (2001). *Core French 5 to 12*. Ministry of Education, BC.

Beeli, G., Esslen, M., & Jäncke, L. (2007). Frequency correlates in grapheme-color synaesthesia. *Psychological Science*, *18*(9), 788-792.

Bernard, J. W. (1986). Messiaen's Synaesthesia, The Correspondence Between Color and Sound Structure in his Music. *Music Perception*, *4*(1), 41-68.

Bialystock, E., & Herman, J. (1999). Does bilingualism matter for early literacy? *Bilingualism: Language and Cognition*, *2*(1), 35-44.

Bialystock, E., Luk, G., & Kwan, E. (2005). Bilingualism, biliteracy, and learning to read: Interactions among languages and writing systems. *Scientific Studies of Reading*, *9*(1), 43-61.

Bialystok, E., & Barac, R. (2012). Emerging bilingualism: dissociating advantages for metalinguistic awareness and executive control. *Cognition*, *122*(1), 67-73. doi:10.1016/j.cognition.2011.08.003

Binet, A., & Philippe, J. (1892). Êtude sur un nouveau cas d'audition colorée. *Revue Philosophique de la France et de l'étranger*, *33*, 461-464.

Bloch, C., . . . Nitsch, C. (2009). The age of second language acquisition determines the variability in activation elicited by narration in three languages in Broca's and Wernicke's area. *Neuropsychologia*, *47*(3), 625-633. doi:10.1016/j.neuropsychologia.2008.11.009

Boles, D. B., & Clifford, J. E. (1989). An upper- and lowercase alphabetic similarity matrix, with derived generation similarity values. *Behavior Research Methods, Instruments, & Computers*, *21*(6), 579-586.

Bor, D., Billington, J., & Baron-Cohen, S. (2007). Savant memory for digits in a case of synaesthesia and Asperger syndrome is related to hyperactivity in the lateral prefrontal Cortex. *Neurocase*, *13*(5-6), 311-319.

Brang, D., Hubbard, E. M., Coulson, S., Huang, M., & Ramachandran, V. S. (2010). Magnetoencephalography reveals early activation of V4 in grapheme-color synesthesia. *Neuroimage*, *53*(1), 268-274.

Brang, D., & Ramachandran, V. S. (2011). Survival of the synesthesia gene: Why do people hear colors and taste words? *PLoS Biology*, *9*(11), e1001205. doi:10.1371/journal.pbio.1001205.g002

Brang, D., Rouw, R., Ramachandran, V. S., & Coulson, S. (2010). *Similar letters, similar hues: Shape-color isomorphism in grapheme-color synaesthesia*. Proceedings from 2010 Meeting of the UK Synaesthesia Association, Brighton, U.K.

Brang, D., Rouw, R., Ramachandran, V. S., & Coulson, S. (2011). Similarly shaped letters evoke similar colors in grapheme-color synesthesia. *Neuropsychologia*, *49*(5), 1355-1358. doi:10.1016/j.Neuropsychologia.2011.01.002

Breslow, L. A., Trafton, J. G., McCurry, J. M., & Ratwani, R. M. (2010). An algorithm for generating color scales for both categorical and ordinal coding. *Color Research and Application*, *35*(1), 18-28. doi:10.1002/col.v35:1

Březinová, L., Havel, J., & Stadlerová, H. (2007). *Živá Abeceda, Učebnice pro 1. ročník základní školy [The Living ABC's: Program in reading for the first year of school]*. Pizeň: Fraus.

Calkins, M. W. (1893). A statistical study of pseudo-chromesthesia and of mental-forms. *American Journal of Psychology*, *5*, 439-464.

Calkins, M. W. (1895). Synaesthesia. *American Journal of Psychology*, *7*, 90-107.

Claparède, E. (1900). Sur l'Audition Colorée. *Revue Philosophique de la France et de l'étranger*, *49*, 515-521.

Clements, G. N., & Hume, E. V. (1995). The internal organization of speech sounds. In J. A. Goldsmith (Ed.), *The Handbook of Phonological Theory* (pp. 245-306). Cambridge, MA: Blackwell.

Cobo-Lewis, A. B., Pearson, P. Z., Eilers, R. E., & Umbel, V. C. (2002). Effects of bilingualism and bilingual education on oral and written Spanish skills: A multifactor study of standardized test outcomes. In D. K. Oller & R. E. Eilers (Eds.), *Language and literacy in bilingual children* (pp. 98-117). Clevedon, UK: Multiligual Matters.

Cohen Kadosh, R., Henik, A., & Walsh, V. (2007). Small is bright and big is dark in synaesthesia. *Current Biology*, *17*(19), R834-R835.

Cohen Kadosh, R., Kadosh, K. C., & Henik, A. (2007). The neuronal correlate of bidirectional synesthesia: A combined event-related potential and functional magnetic resonance Imaging study. *Journal of Cognitive Neuropsychology*, *19*(12), 2050-2059.

Cohen Kadosh, R., & Terhune, D. B. (2012). Redefining synaesthesia? *British Journal of Psychology*, *103*, 20-23.

Colizoli, O., Murre, J. M. J., Rouw, R., Karniel, A., & Witthoft, N. (2012). Pseudo-Synesthesia through Reading Books with Colored Letters. *PLoS One*, *7*(6), 1-10.

Curtis, C. (1998). Letter-color synaesthesia. *OtherThings.com*. Retrieved April 12, 2013, from otherthings.com/uw/syn

Cytowic, R. E. (1988). Tasting colors, smelling sounds - Neurological clues to a confounding condition. *The Sciences*, 32-37.

Cytowic, R. E. (1989a). *Synesthesia: A Union of the Senses* (First ed.). New York: Springer-Verlag.

Cytowic, R. E. (1989b). Synesthesia and mapping of subjective sensory dimensions. *Neurology*, *39*(6), 849-850.

Cytowic, R. E. (1993). *The Man Who Tasted Shapes*. New York: G.P. Putnam's Sons.

Cytowic, R. E. (1997). Synaesthesia: phenomenology and neuropsychology - a review of current knowledge. In S. Baron-Cohen & J. E. Harrison(pp. 17-39). Oxford: Basil Blackwell.

Cytowic, R. E. (2002). *Synesthesia: A Union of the Senses* (Second ed.). Cambridge Massachusetts: The MIT Press.

Cytowic, R. E., & Eagleman, D. M. (2009). *Wednesday is Indigo Blue: Discovering the Brain of Synesthesia*. Boston, MA: MIT Press.

Cytowic, R. E., & Wood, F. B. (1982a). Synesthesia I. A review of major theories and their brain basis. *Brain and Cognition*, *1*, 23-35.

Cytowic, R. E., & Wood, F. B. (1982b). Synesthesia II. Psychophysical relations in the synesthesia of geometrically shaped taste and colored hearing. *Brain and Cognition*, *1*, 36-49.

Dailey, A., Martindale, C., & Borkum, J. (1997). Creativity, synesthesia, and physiognomic perception. *Creativity Research Journal*, *10*(1), 1-8.

Day, S. A. (2005). Some demographic and socio-cultural aspects of synesthesia. In L. C. Robertson & N. Sagiv(pp. 11-33). Oxford: Oxford University Press.

Dindia, K., & Allen, M. (1992). Sex differences in self-disclosure: A meta-analysis. *Psychological Bulletin*, *112*(1), 106-124.

Dixon, M. J., Smilek, D., Cudahy, C., & Merikle, P. M. (2000). Five plus two equals yellow: Mental arithmetic in people with synaesthesia is not coloured by visual experience. *Nature*, *406*(679), 365-365.

Dixon, M. J., Smilek, D., & Merikle, P. M. (2004). Not all synaesthetes are created equal: Projector versus associator synaesthetes. *Cognitive, Affective, & Behavioral Neuroscience*, *4*, 335-343.

Domino, G. (1989). Synesthesia and creativity in fine arts students: An empirical look. *Creativity Research Journal*, *2*, 17-29.

Eagleman, D. M. (2010). *What has large scale analysis taught us?* Proceedings from 2010 Meeting of the UK Synaesthesia Association, Brighton, U.K.

Eagleman, D. M. (2012). Synaesthesia in its protean guises. *British Journal of Psychology*, *103*, 16-19.

Eagleman, D. M., Kagan, A. D., Nelson, S. S., Sagaram, D., & Sarma, A. K. (2007). A standardized test battery for the study of synesthesia. *Journal of Neuroscience Methods*, *159*(1), 139-145.

Ellis, N. C., . . . Petalas, M. (2004). The effects of orthographic depth on learning to read alphabetic, syllabic, and logographic scripts. *Reading Research Quarterly*, *39*(4), 438-468.

Flournoy, T. (1892). L'audition colorée. *Archives des Sciences Physiques et Naturelles*, *28*, 505-508.

Frisch, S. A. (1996). *Similarity and Frequency in Phonology.* PhD. Northwestern University, Evanston, Il.

Frisch, S. A., Pierrehumbert, J. B., & Broe, M. B. (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory*, *22*, 179-228.

Galton, F. (1883). Inquiries into Human Faculty and Its Developmment.

Gheorghiu, E., & Kingdom, F. A. A. (2006). Luminace-contrast properties of contour-shape processing revealed through the shape-frequency after-effect. *Vision Research*, *46*(21), 3603-3615.

Gheorghiu, E., & Kingdom, F. A. A. (2007). Chromatic tuning of contour-shape mechanisms revealed through the shape-frequency and shape-amplitude after-effects. *Vision Research*, *47*(14), 1935-1949. doi:10.1016/j.visres.2007.03.010

Ghirardelli, T. G., Mills, C. B., Zilioli, M. K. C., Bailey, L. P., & Kretschmar, P. K. (2010). Synesthesia affects verification of simple arithmetic equations. *Journal of General Psychology*, *137*(2), 175-189.

Gibson, E. J. (1969). *Principles of Perceptual Learning and Development*. New Jersey: Prentice-Hall, Inc.

Gilmore, G. C., Hersh, H., Caramazza, A., & Griffin, J. (1979). Multidimensional letter similarity derived from recognition errors. *Perception & Psychophysics*, *25*(5), 425-431.

Green, J. A. K., & Goswami, U. (2008). Synesthesia and number cognition in children. *Cognition*, *106*(1), 463-473.

Grossenbacher, P. G. (1997). Perception and sensory information in synesthetic experience. In S. Baron-Cohen & J. E. Harrison(pp. 148-172). Oxford: Basil Blackwell.

Grossenbacher, P. G., & Lovelace, C. T. (2001). Mechanisms of synesthesia: cognitive and physiological constraints. *Trends in Cognitive Sciences*, *5*(1), 36-41.

Gupta, S. M., Geyer, L. H., & Maalouf, J. A. (1983). Effect of font and medium on recognition/confusion. *SIGCHI conference on Human Factors in Computing Systems*, 144-149.

Hancock, P. (2006). Monozygotic twins' colour-number association: A case study. *Cortex*, *42*(2), 147-150.

Hänggi, J., Beeli, G., Oechslin, M. S., & Jancke, L. (2008). The multiple synaesthete E.S.: neuroanatomical basis of interval-taste and tone-colour synaesthesia. *Neuroimage*, *43*(2), 192-203. doi:10.1016/j.Neuroimage.2008.07.018

Hansen, T., & Gegenfurtner, K. R. (2009). Independence of color and luminance edges in natural scenes. *Visual Neuroscience*, *26*(1), 35-49.

Head, P. D. T. (2006). Synaesthesia: Pitch-colour isomorphism in RGB-space? *Cortex*, *42*(2), 164-174.

Hubbard, E. M., Arman, A. C., Ramachandran, V. S., & Boynton, G. M. (2005). Individual differences among grapheme-color synesthetes: Brain-behavior correlations. *Neuron*, *45*(6), 975-985.

Hubbard, E. M., Brang, D., & Ramachandran, V. S. (2011). The cross-activation theory at 10. *Journal of Neuropsychology*, *5*(2), 152-177. doi:10.1111/j.1748-6653.2011.02014.x

Ione, A. (2004). Klee and Kandinsky: Polyphonic painting, chromatic chords, and synaesthesia. *Journal of Consciousness Studies*, *11*(3-4), 148-158.

Jäncke, L., Rogenmoser, L., Meyer, M., & Elmer, S. (2012). Pre-attentive modulation of brain responses to tones in coloured-hearing synesthetes. *BMC Neuroscience*, *13*(1), 1-15. doi:10.1186/1471-2202-13-151

Jewanski, J., Day, S. A., & Ward, J. (2009). A Colorful Albino: The First Documented Case of Synaesthesia, by Georg Tobias Ludwig Sachs in 1812. *Journal of the History of the Neurosciences*, *18*(3), 293-303.

Jürgens, U. M., & Nikolic, D. (2012). Ideasthesia: Conceptual processes assign similar colours to similar shapes. *Translational Neuroscience*, *3*(1), 22-27.

Justice, L. M., Pence, K., Bowles, R. B., & Wiggins, A. (2006). An investigation of four hypotheses concerning the order by which 4-year-old children learn the alphabet letters. *Early Childhood Research Quarterly*, *21*(3), 374-389. doi:10.1016/j.ecresq.2006.07.010

Kingdom, F. A. A., Beauce, C., & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception*, *33*(8), 907-914.

Kingdom, F. A. A., & Kasrai, R. (2006). Colour unmasks dark targets in complex displays. *Vision Research*, *46*(6-7), 814-822. doi:10.1016/j.visres.2005.08.018

Kovács, Á. M., & Mehler, J. (2009). Cognitive gains in 7-month-old bilingual infants. *PNAS*, *106*(16), 6556-6560.

Králik, J. (1983). Statistika českých grafémů s využitím moderní výpočetní techniky [Statistics of the Czech graphemes with the aid of modern computational techniques]. *Slove a Slovesnost [Word and Literature]*, *44*, 295-304.

Kusnir, F., & Thut, G. (2012). Formation of automatic letter–colour associations in non-synaesthetes through likelihood manipulation of letter–colour pairings. *Neuropsychologia*, *50*(14), 3641-3652. doi:10.1016/j.neuropsychologia.2012.09.032

Ladová, A., Holas, M., & Staudková, H. (2011). *Živá Abeceda, Učebnice pro 1. ročník základních škol pro vzdělávací obor Český jazyk a literature [The Living ABC's: Program in reading for the first year of school with Czech rhymes and literature]*. Praha: Alter.

Laeng, B. (2009). Searching through synaesthetic colors. *Attention, Perception, & Psychophysics*, *71*(7), 1461-1467.

Laeng, B., Hugdahl, K., & Specht, K. (2011). The neural correlate of colour distances revealed with competing synaesthetic and real colours. *Cortex*, *47*(3), 320-331. doi:10.1016/j.Cortex.2009.09.004

Lewand, R. (2000). *Cryptological Mathematics*. Washington: The Mathematical Association of America.

Liebe, S., Fischer, E., Logothetis, N., & Rainer, G. (2009). Color and shape interactions in the recognition of natural scenes by human and monkey observers. *Journal of Vision*, *9*(5), 1-16.

Luria, A. R. (1968). *The Mind of a Mnemonist*. New York: Basic Books.

MacLeod, C. M., & Dunbar, K. (1988). Training and Stroop-like interference: Evidence for a continuum of automaticity. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*(1), 126-135.

Maddox, W. T., Filoteo, J. V., Hejl, K. D., & Ing, A. D. (2004). Category number impacts rule-based but not information-integration category learning: Further evidence for dissociable category-learning systems. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *30*(1), 1-19. doi:10.1037/0278-7393.30.1.227

Marks, L. E. (1975). On colored-hearing synesthesia: Cross-modal translations of sensory dimensions. *Psychological Bulletin*, *82*(3), 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. MacNeilage, & J. Morton(pp. 109-124). New York: Springer. Retrieved from http://direct.bl.uk/research/26/64/EN004127304.html?source=googlescholar&oi=docdel

MedCalc Software. (2013). Relative Risk. *Medcalc*. Retrieved April 14, 2013, from http://www.medcalc.org/calc/relative_risk.php

Melichárková, I., Štěpán, L., & Švecová, L. (2008). *Slabikář [The Syllable Book]*. Liberec: Dialog.

Mikulenková, H., & Mladý, R. (2004). *Příručka k vyučování čtení, psaní a literature v prvním ročníku základní školy [Teaching Reading to 6 year olds, with literature for the first year of school]*. Olomouc: Prodos.

Mikulenková, H., Mladý, R., & Forman, M. (2004). *Slabikář [The Syllable Book]*. Olomouc: Prodos.

Mills, C. B., Boteler, E. H., & Oliver, G. K. (1999). Digit synaesthesia: A case study using a Stroop-type test. *Cognitive Neuropsychology*, *16*(2), 181-191.

Mills, C. B., Viguers, M. L., Edelson, S. K., Thomas, A. T., Simon-Dack, S. L., & Innis, J. A. (2002). The color of two alphabets for a multilingual synesthete. *Perception*, *31*(11), 1371-1394. doi:10.1068/p3429

Nabokov, V. (1989). *Speak, Memory: An Autobiography Revisited*. New York: Random House, Inc.

Nagai, T., & Uchikawa, K. (2009). Different hue coding underlying figure segregation and region detection tasks. *Journal of Vision*, *9*(9), 1-19.

Niccolai, V. (2012). Modality and Variability of Synesthetic Experience. *The American Journal of Psychology*, *125*(1), 81-94. doi:10.5406/amerjpsyc.125.1.0081

Nikolic, D., Jurgens, U. M., Rothen, N., Meier, B., & Mroczko, A. (2011). Swimming-style synesthesia. *Cortex*, *47*(7), 874-879. doi:10.1016/j.Cortex.2011.02.008

Nováčková, O. (2010). *Živá Abeceda, Učebnice vytvořená v souladu s RVP ZV [The Living ABC's: Program for reading and writing for the first year at school]*. Brno: Nova Skola.

Novich, S., Cheng, S., & Eagleman, D. M. (2011). Is synaesthesia one condition or many? A large-scale analysis reveals subgroups. *Journal of Neuropsychology*, *5*(2), 353-371. doi:10.1111/j.1748-6653.2011.02015.x

Nunn, J. A., . . . Gray, J. A. (2002). Functional magnetic resonance imaging of synesthesia: activation of V4/V8 by spoken words. *Nature Neuroscience*, *5*(4), 371-375.

Ortiz-Mantilla, S., Choudhury, N., Alvarez, B., & Benasich, A. A. (2010). Involuntary switching of attention mediates differences in event-related responses to complex tones between early and late Spanish-English bilinguals. *Brain Res*, *1362*, 78-92. doi:10.1016/j.brainres.2010.09.031

Pariyadath, V., Plitt, M. H., Churchill, S. J., & Eagleman, D. M. (2012). Why overlearned sequences are special: distinct neural networks for ordinal sequences. *Frontiers in Human Neuroscience*, *6*, 1-9. doi:10.3389/fnhum.2012.00328/abstract

Paulesu, E., . . . Frith, C. D. (1995). The physiology of coloured hearing: A PET activation study of colour-word synaesthesia. *Brain*, *118*, 661-676.

Pautzke, R. (2010). *'Making sense' messing around with black boxes: Synaesthesia and learning - How to master the Theremin without notes*. Proceedings from 2010 Meeting of the UK Synaesthesia Association, Brighton, UK.

Peacock, K. (1985). Synaesthetic perception: Alexander Scriabin's color hearing. *Music Perception*, *2*(4), 483-506.

Podgorny, P., & Garner, W. R. (1979). Reaction time as a measure of inter- and intraobject visual similarity: Letters of the alphabet. *Perception & Psychophysics*, *26*(1), 37-52.

Potůčková, J. (2010). *Živá Abeceda pro 1. ročník základní školy [The Living ABC's: Program #1 for the first year of school]*. Brno: Studio 11.

Radvansky, G. A., Gibson, B. S., & McNerney, M. W. (2011). Synesthesia and memory: Color congruency, von Restorff, and false memory effects. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *37*(1), 219-229.

Ramachandran, V. S., & Hubbard, E. M. (2001a). Synaesthesia--A Window into Perception, Thought and Language. *Journal of Consciousness Studies*, *8*(12), 3-34.

Ramachandran, V. S., & Hubbard, E. M. (2001b). Psychophysical investigations into the neural basis of synaesthesia. *Proceedings of the Royal Society of London Series B - Biological Sciences*, *268*(1470), 979-983.

Ramachandran, V. S., & Rogers-Ramachandran, D. (1996). Synaesthesia in phantom limbs induced with mirrors. *Proceedings of the Royal Society of London Series B - Biological Sciences*, *263*(1369), 377-386.

Reber, A. S. (1993). *Implicit Learning and Tacit Knowledge: An Essay on the Cognitive Unconscious*. Oxford: Oxford University Press.

Rich, A. N., Bradshaw, J. L., & Mattingley, J. B. (2005). A systematic, large-scale study of synaesthesia: Implications for the role of early experience in lexical-colour associations. *Cognition*, *98*, 53-84.

Rich, A. N., Williams, M. A., Puce, A., Syngeniotis, A., Howard, M., McGlone, F., & Mattingley, J. B. (2006). Neural correlates of imagined and synaesthetic colours. *Neuropsychologia*, *44*(1), 2918-2925.

Rose, K. B. (1909). Some statistics on synaesthesia. *American Journal of Psychology*, *20*(3), 447.

Rothen, N., & Meier, B. (2009). Do synesthetes have a general advantage in visual search and episodic memory? A case for group studies. *PLoS One*, *4*(4), e5037. doi:10.1371/journal.pone.0005037

Rothen, N., & Meier, B. (2010a). Grapheme-colour synaesthesia yields an ordinary rather than extraordinary memory advantage: Evidence from a group study. *Memory*, *18*(3), 258-264.

Rothen, N., & Meier, B. (2010b). Higher prevalence of synaesthesia in art students. *Perception*, *39*(5), 718-720.

Rothen, N., Meier, B., & Ward, J. (2012). Enhanced memory ability: Insights from synaesthesia. *Neuroscience and Biobehavioral Reviews*, *36*(8), 1952-1963. doi:10.1016/j.neubiorev.2012.05.004

Rouw, R., & Scholte, H. S. (2007). Increased structural connectivity in grapheme-color synesthesia. *Nature Neuroscience*, *10*(6), 792-797. Retrieved from http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?cmd=Retrieve&db=PubMed&dopt=Citation&list_uids=17515901

Rouw, R., & Scholte, H. S. (2010). Neural Basis of Individual Differences in Synesthetic Experiences. *Journal of Neuroscience*, *30*(18), 6205-6213.

Rouw, R., Scholte, H. S., & Colizoli, O. (2011). Brain areas involved in synaesthesia: A review. *Journal of Neuropsychology*, *5*(2), 214-242. doi:10.1111/j.1748-6653.2011.02006.x

Saenz, M., & Koch, C. (2008). The sound of change: visually-induced auditory synesthesia. *Current Biology*, *18*(15), R650-R651. doi:10.1016/j.cub.2008.06.014

Sagiv, N., Simner, J., Collins, J., Butterworth, B., & Ward, J. (2006). What is the relationship between synaesthesia and visuo-spatial number forms? *Cognition*, *101*, 114-128.

Seaberg, M. (2012). Synesthetes: “People of the future”. *Psychology Today*. Retrieved March 16, 2013, from http://www.psychologytoday.com/blog/tasting-the-universe/201203/synesthetes-people-the-future

Seymour, P. H. K., Aro, M., & Erskine, J. M. (2003). Foundation literacy acquisition in European orthographies. *British Journal of Psychology*, *94*, 143-174.

Shimono, K., Shiori, S., & Yaguchi, H. (2009). Psychophysical evidence for a purely binocular color system. *Vision Research*, *49*(2), 202-210.

Simner, J. (2012a). Defining synaesthesia. *British Journal of Psychology*, *103*, 1-15.

Simner, J. (2012b). Defining synaesthesia: A response to two excellent commentaries. *British Journal of Psychology*, *103*, 24-27.

Simner, J., Glover, L., & Mowat, A. (2006). Linguistic determinants of word colouring in grapheme-colour synaesthesia. *Cortex*, *42*(2), 281-289.

Simner, J., Harrold, J., Creed, H., Monro, L., & Foulkes, L. (2009). Early detection of markers for synaesthesia in childhood populations. *Brain*, *132*(Pt 1), 57-64. doi:10.1093/brain/awn292

Simner, J., Hung, W. Y., & Shillcock, R. (2011). Synaesthesia in a logographic language: The colouring of Chinese characters and Pinyin/Bopomo spellings. *Consciousness and Cognition*, *20*(4), 1376-1392.

Simner, J., Mayo, N., & Spiller, M. J. (2009). A foundation for savantism? Visuo-spatial synaesthetes present with cognitive benefits. *Cortex*, *45*(10), 1246-1260.

Simner, J., . . . Ward, J. (2006). Synaesthesia: The prevalence of atypical cross-modal experiences. *Perception*, *35*, 1024-1033.

Simner, J., & Ward, J. (2008). Synaesthesia, color terms, and color space - Color claims came from color names in Beeli, Esslen, and Jancke (2007). *Psychological Science*, *19*(4), 412-414.

Simner, J., Ward, J., Lanz, M., Jansari, A., Noonan, K., Glover, L., & Oakley, D. A. (2005). Non-random associations of graphemes to colours in synaesthetic and non-synaesthetic populations. *Cognitive Neuropsychology*, *22*, 1069-1085.

Skelton, R., Ludwig, C., & Mohr, C. (2009). A novel, illustrated questionnaire to distinguish projector and associator synaesthetes. *Cortex*, *45*(6), 721-729.

Smilek, D., Carriere, J. S. A., Dixon, M. J., & Merikle, P. M. (2007). Grapheme frequency and color luminance in grapheme-color synaesthesia. *Psychological Science*, *18*(9), 793-795.

Smilek, D., Dixon, M. J., Cudahy, C., & Merikle, P. M. (2001). Synaesthetic photisms influence visual perception. *Journal of Cognitive Neuroscience*, *13*, 930-936.

Smilek, D., Dixon, M. J., Cudahy, C., & Merikle, P. M. (2002a). Concept driven color experiences in digit-color synesthesia. *Brain and Cognition*, *48*, 570-573.

Smilek, D., Dixon, M. J., Cudahy, C., & Merikle, P. M. (2002b). Synesthetic color experiences influence memory. *Psychological Science*, *13*, 548-552.

Spector, F., & Maurer, D. (2008). The colour of Os: Naturally biased associations between shape and colour. *Perception*, *37*(6), 841-847. doi:10.1068/p5830

Spector, F., & Maurer, D. (2011). The colors of the alphabet: Naturally-biased associations between shape and color. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(2), 484-495. doi:10.1037/a0021437

Tomson, S. N., . . . Eagleman, D. M. (2011). The genetics of colored sequence synesthesia: Suggestive evidence of linkage to 16q and genetic heterogeneity for the condition. *Behavioural Brain Research*, *223*(1), 48-52. doi:10.1016/j.bbr.2011.03.071

Ulich, E. (1957). Synästhesie und geschlecht - Empirische Untersuchungen über die geschlechtsspezifische Häufigkeit des Auftretens synästhetischer Phänomene [Synesthesia and gender - empirical studies on the gender frequency of occurrence of synaesthetic phenomena]. *Zeitschrift Fur Experimentelle und Angewandte Psycholgie*, *4*, 31-52.

Ward, J., Li, R., Salih, S., & Sagiv, N. (2007). Varieties of grapheme-colour synaesthesia: A new theory of phenomenological and behavioural differences. *Consciousness and Cognition*, *16*(4), 913-931.

Ward, J., & Simner, J. (2003). Lexical-gustatory synaesthesia: Linguistic and conceptual factors. *Cognition*, *89*(3), 237-261.

Ward, J., & Simner, J. (2005). Is synaesthesia an X-linked dominant trait with lethality in males? *Perception*, *34*(5), 611-623.

Ward, J., Simner, J., & Auyeung, V. (2005). A comparison of lexical-gustatory and grapheme-colour synaesthesia. *Cognitive Neuropsychology*, *22*, 28-41.

Ward, J., Thompson-Lake, D., Ely, R., & Kaminski, F. (2008). Synaesthesia, creativity and art: What is the link? *British Journal of Psychology*, *99*, 127-141.

Ward, J., Tsakanikos, E., & Bray, A. (2006). Synaesthesia for Reading and Playing Musical Notes. *Neurocase*, *12*, 27-34.

Watson, M. R., Akins, K., & Crawford, L. (2010). The developmental learning hypothesis of synaesthesia - A Summary. *Studie z aplikované lingvistiky/Studies in Applied Linguistics*, *1*.

Watson, M. R., Akins, K. A., & Enns, J. T. (2012). Second-order mappings in grapheme-color synesthesia. *Psychonomic Bulletin and Review*, *19*(2), 211-217. doi:10.3758/s13423-011-0208-4

Watson, M. R., Blair, M. R., Kozik, P., Akins, K. A., & Enns, J. T. (2012). Grapheme-color synaesthesia benefits rule-based category learning. *Consciousness and Cognition*, *21*, 1533-1540. doi:10.1016/j.concog.2012.06.004

Weiss, P. H., & Fink, G. R. (2009). Grapheme-colour synaesthetes show increased grey matter volumes of parietal and fusiform cortex. *Brain*, *132*(Pt 1), 65-70. doi:10.1093/brain/awn304

Weiss, P. H., Zilles, K., & Fink, G. R. (2005). When visual perception causes feeling: Enhanced cross-modal processing in grapheme-color synesthesia. *Neuroimage*, *28*(4), 859-868. doi:10.1016/j.neuroimage.2005.06.052

Werker, J. F., & Byers-Heinlein, K. (2008). Bilingualism in infancy: first steps in perception and comprehension. *Trends in Cognitive Sciences*, *12*(4), 144-151.

Wheeler, R. H., & Cutsforth, T. D. (1921). The role of synaesthesia in learning. *Journal of Experimental Psychology*, *4*, 448-468.

Wheeler, R. H., & Cutsforth, T. D. (1922). Synaesthesia and meaning. *American Journal of Psychology*, *33*, 361-384.

Wheeler, R. H., & Cutsforth, T. D. (1925). Synæsthesia in the development of the concept. *Journal of Experimental Psychology*, *8*, 149-159.

Witthoft, N., & Winawer, J. (2006). Synesthetic colors determined by having colored refrigerator magnets in childhood. *Cortex*, *42*(2), 175-183.

Witthoft, N., & Winawer, J. (2013). Learning, memory, and synesthesia. *Psychological Science*. doi:10.1177/0956797612452573

Yaro, C., & Ward, J. (2007). Searching for Shereshevskii: What is superior about the memory of synaesthetes? *The Quarterly Journal of Experimental Psychology*, *60*(5), 681-695.

Žáček, J., & Zmatlíková, H. (2010). *Slabikář [The Syllable Book]*. Praha: Prodos.

Appendix 1 - Synaesthesia Survey

Email: ____

Gender:
____ Male
____ Female

Handedness:
____ Left-handed
____ Right-handed
____ Both

First language (the one you learned to understand first)(if you learned two languages from birth, please name both languages): _________________________________

Have you learned any other languages? What age were you when you learned them?

Language: ____________________________ Age: ________
Language: ____________________________ Age: ________
Language: ____________________________ Age: ________

When you see, hear or think about certain *letters* or *numbers*, do you see or feel any *colours*? (Example: There is something yellow about the letter **G**.)
____ No, I do not have experiences like this.
____ Yes, I have experiences like this with *letters*.
____ Yes, I have experiences like this with *numbers*.

When you see, hear or think about the *names of the days of the week* or *months*, do you see or feel any *colours*? (Example: The name "Monday" is purple.)
____ No, I do not have experiences like this.
____ Yes, I have experiences like this.

When you hear certain *sounds*, do you see or feel any *colours*? (Examples: Car horns are blue. The musical note C-sharp feels dark green. The sound of a guitar seems pink.)
____ No, I do not have experiences like this.
____ Yes, I have experiences like this.

Do any *words* seem to have *tastes or smells*? (Example: The word "elephant" tastes like strawberries.)
____ No, I do not have experiences like this.
____ Yes, I have experiences like this.

Do any *numbers, hours, days of the week, or months* seem to have a *location in space around you*? (Example: September is always three feet to my left.)
____ No, I do not have experiences like this.
____ Yes, I have experiences like this.

Turn Over →

Do you think of any *numbers* or *letters* as having a *personality, gender, or age*? (Example: 4 is a friendly, older lady.)
____ No, I do not have experiences like this.
____ Yes, I have experiences like this.

These questions have described some different kinds of synaesthesia, but there are other kinds, too. Do you think you might have another kind that we did not ask about?
____ No.
____ Yes. Describe:____

Do you remember telling yourself stories with *letters or numbers as characters* when you were a child?
____ No, I do not remember telling myself stories like this.
____ Yes, I remember telling myself stories like this.

Did you learn to read *before you went to kindergarten* (or whatever kind of school you went to at *age 5*)?
_____I do not know.
_____Yes, I learned to read before I went to kindergarten.
_____No, I learned to read after I went to kindergarten.

As a child, were you ever given *extra help with reading, writing or spelling*? For example, did you ever have one-on-one lessons with a special teacher?
____I do not remember.
____ No, I was not.
____ Yes, I was.

As a child, were you ever told you had *dyslexia* or some other *problem with reading*?
____I do not remember.
____ No, I was not.
____ Yes, I was.

Do you feel that *reading or spelling is difficult* for you now?
____ I am not sure.
____ No, these things are not difficult.
____ Yes, these things are difficult.

Appendix 2 - Associations between reported synaesthesia types

	Letter	Number	Wkdy/Mth	Sound	Word-Taste	Number Fm	Other
Letter-Colour		3.00	3.54	2.70	2.26	2.34	2.68
Number-Colour	4.79		4.45	3.86	2.33	2.27	2.55
Weekday/ Month-Colour	4.11	4.91		4.24	2.83	2.95	2.50
Sound-Colour	2.07	2.33	1.96		3.65	3.40	3.59
Word-Taste	1.63	1.62	1.60	2.59		2.99	2.10
Number Forms	1.57	1.58	1.45	1.66	1.54		2.04
Other Types	2.56	2.69	1.79	2.27	1.46	1.99	

Relative rates of simultaneous occurrence of different forms of reported synaesthesia (how much more likely someone who reports synaesthesia of Type A is to also report synaesthesia of Type B than someone who does not report synaesthesia of Type A). The upper triangle is taken from the SFU sample, the lower from the CU sample. *P*-values are all <.001 after Bonferroni correction.

Appendix 3 - Associations between confirmed synaesthesia types

	Letter	Number	Weekday	Month	Chord	Scale	Instrument
Letter-Colour		66.53***	25.75***	36.96***	44.35	79.84*	133.06
Number-Colour	65.98***		52.86***	49.22***	0.00	41.65	41.65
Weekday-Colour	47.07***	49.70***		93.51***	37.11	66.79	0.00
Month-Colour	50.18***	43.66***	52.11***		89.35*	148.91*	89.35
Chord-Colour	0.00	35.26*	0.00	54.31		0.00	346.00
Scale-Colour	25.94**	23.90**	34.96	41.19	111.55		172.83
Instrument-Colour	48.15	27.23	13.72	13.56	0.00	79.71**	

Relative rates of simultaneous occurrence of different forms of confirmed synaesthesia (how much more likely someone who has synaesthesia of Type A is to also have synaesthesia of Type B than someone who does not have synaesthesia of Type A). The upper triangle is taken from the SFU sample, the lower from the CU sample. *P*-values are all Bonferroni-corrected, and > .05 except: *** $p < .001$, ** $p < .01$, * $p < .05$