The Face Inversion Effect: Investigating the role of perceptual learning, facial specificity, and holistic processing.

Submitted by Siobhan McCourt, to the University of Exeter as a thesis for the degree of Doctor of Philosophy in Psychology February 2024.

I certify that all material in this thesis which is not my own work has been identified and that any material that has previously been submitted and approved for the award of a degree by this or any other University has been acknowledged.

ABSTRACT

This thesis investigates the perceptual processes underlying face recognition and the face inversion effect, exploring whether there is evidence for facial specificity in this effect and the specific types of information that produce it. Results will be presented from behavioural studies using a variety of manipulated face stimuli and studies involving transcranial direct current stimulation (tDCS) using face stimuli and prototype-defined checkerboard stimuli, which have previously been used to demonstrate the role of perceptual expertise in the face inversion effect (McLaren, 1997; Civile, Zhao, et al., 2014; Civile, Verbruggen, et al., 2016). Chapter 1 outlines the previous literature and background theories underpinning the face recognition research. Chapter 2 directly compares the effect of tDCS on the inversion effect for faces and checkerboards and tests a new active control condition. The disparity in the remaining inversion effect for faces and checkerboards under tDCS has led to the suggestion that there may be an additional, potentially face-specific component contributing to the inversion effect for faces together with perceptual learning. The findings here offer some support for this idea and also indicate based on the active control comparison that it is the specific Fp3-Fp2 montage that produces this tDCS induced effect. Chapter 3 begins a series of experiments exploring the idea that holistic processing (indexed by face contour) may be part of this additional component. Scrambled faces were used on the basis that they have been shown to result in a robust inversion effect despite complete disruption to configural information (Civile et al., 2014), and were subject to a contour manipulation to assess the impact of this on the inversion effect. Results showed that disruption to the contour information in addition to scrambling was sufficient to reduce the inversion effect. Chapters 4 and 5 extend this contour manipulation to normal faces

and New Thatcherised faces to further explore the impact on the inversion effect. Results from these studies are somewhat mixed with some showing that contour manipulation reduces overall performance only, and others indicating that it impacts the inversion effect. Chapter 6 aims to investigate how tDCS stimulation is able to produce the effect on face recognition that it has been shown to and does so by utilising the typical anodal Fp3-Fp2 montage and then reversing the polarity to deliver cathodal stimulation. This reversal was shown to also reverse the behavioural effects, with anodal stimulation resulting in a reduction to the inversion effect and subsequently delivered cathodal stimulation increasing it again. Chapter 7 summarises the experimental findings and discusses the implications in terms of the wider literature as well as offering suggestions for future research.

Acknowledgements

I would firstly like to thank my supervisors Dr. Ciro Civile and Professor Ian McLaren, their constant support has taken me from my MSc through to the completion of my PhD. Ciro's enthusiasm and obvious passion for the subject inspired my own initial interest in face recognition and neuroscientific techniques and I thank him for sharing this with me and supporting my learning of the systems used to run tDCS experiments. I am grateful to Ian for the insights and suggestions he provided that allowed me get the best out of my work and the advice he offered in terms of my wider career. I am aware that the time and dedication given to my progress by both supervisors is a rare thing and I consider myself incredibly lucky to have had them in my corner throughout my time at Exeter.

I would also like to express my gratitude to Rossy McLaren not only for the invaluable help she gave in contributing to data collection but also for being a consistently positive presence in the lab and fostering the personal connections that served to strengthen the lab as a group.

The support and company of my office mates Toby, Brontë, Kathryn, Emika and Salwa has also been greatly appreciated, knowing that I was not alone in this process was a great comfort and seeing many of them succeed ahead of me reassured me that it was possible.

My parents have been my unwavering and unconditional support system from my first foray into education at age 4 through to this point, celebrating my every success no matter how small along the way. They always taught me to believe that I could achieve anything I wanted to and gave me the freedom to choose what that would be. I have them to thank also for giving me my little sister Lolly who would grow into my best friend and greatest cheerleader.

My friends Katie, Adam and Tash have supported me throughout my academic career from my primary school days to my decision to pursue a PhD, without their help I would not have reached this point. More importantly, they are fantastic friends with whom I share many of my most treasured memories.

Katie, Elliot, Kirsty and Emily make up the greatest group of girls I could have hoped to meet and their friendship has sustained me through the hardest times, irrespective of the physical distance between us. Each one of them intelligent and successful in their own right, they have inspired me to dream big and pushed me to succeed.

I dedicate this thesis to my late Nana who was such a proud and vocal champion of my abilities and who I so wish could have been here to see it completed.

List of Contents

List of Figures

faces with either normal or blurred contour. The y-axis shows d’. Error bars represent s.e.m.

Figure 11. Examples of the stimuli used in Experiment 3b. Panel A shows faces with normal contour and panel B shows faces with manipulated contour.

Figure 12. Graph reports the results from Experiment 3b (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, faces with either normal or manipulated contour. The y-axis shows d’. Error bars represent s.e.m.

Figure 13. Graph reports the results from Experiments 3a and 3b combined (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, faces with either normal or blurred/manipulated contour. The y-axis shows d’. Error bars represent s.e.m.

Figure 14. Examples of the stimuli used in Experiment 4. Categories specify the configuration of the internal features of the face. Panel A shows faces with normal contour and panel B shows faces with manipulated contour.

Figure 15. Graph reports the results from Experiment 4 (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, New Thatcherised faces with either normal or manipulated contour. The y-axis shows d’. Error bars represent s.e.m.

Figure 16. Examples of stimuli used in Experiment 5a and 5b. Upright and inverted faces with the hair cropped, standardised to greyscale on a black background

Figure 17. A representation of the tDCS montages used in Experiments 5a. In the study phase the tDCS montage used was always anodal Fp3-Fp2 and in the recognition phase the tDCS montage used was cathodal Fp3-Fp2 in the anodal-cathodal condition and sham Fp3-Fp2 in the anodal-sham condition

Figure 18: Graph reports the results from Experiment 5a. The x-axis shows the tDCS conditions, anodal-cathodal and anodal-sham. Both conditions have active stimulation during the first half. The y-axis shows d’. Error bars represent s.e.m.

Figure 19. A representation of the tDCS montages used in Experiments 5b. Panel a. shows the set-up in the anodal-cathodal condition, in the study phase the tDCS montage used is anodal Fp3-Fp2 and in the recognition phase the tDCS montage used is cathodal Fp3-Fp2. Panel b. shows the set-up in the anodal-sham condition, in the study phase the tDCS montage used is anodal Fp3-Fp2 and in the recognition phase the tDCS montage used is sham Fp3-Fp2. Panel c. shows the set-up in the sham-sham condition, in the study phase the tDCS montage used with either sham Fp3-Fp2 or sham Fp2-Fp3 and in the recognition phase the tDCS montage used was also either sham Fp3-Fp2 or sham Fp2-Fp3.

Figure 20: Graph reports the results from Experiment 5b. The x-axis shows the tDCS conditions, sham-sham, anodal-cathodal, and anodal-sham. The y-axis shows d'. Error bars represent s.e.m.

Chapter 1: Literature Review

1.1 The Face Inversion Effect: Specificity vs. Expertise

Face recognition is a skill fundamental to the human experience and as such is one that we are highly adept at performing very quickly. It is easy therefore to underappreciate how difficult a skill it actually is; faces share many similar characteristics with one another, with the overall shape and configuration shared between almost every exemplar. They are also regularly viewed from a variety of angles, distances, and lighting conditions. While we perform this skill with relative ease every day with faces in their usual, upright orientation, the true challenge of recognising these stimuli is revealed when they are presented in the inverted orientation (upside-down). This deficit in recognition ability when stimuli are inverted is referred to as the inversion effect and has been shown to be greater for faces than for objects (e.g., Yin, 1969; Valentine & Bruce, 1986). In early research this was interpreted as evidence of a face-specific recognition/processing mechanism and thus it was dubbed the Face Inversion Effect. Whether this interpretation continues to be supported or if findings can be better explained by expertise with faces has been an area of great debate throughout the literature and fundamentally underpins the question *Are faces special?*

1.2.1 The Specificity account: Yin's, (1969) findings and support for a specificity interpretation

The existence of an inversion effect for faces had initially been shown in both memory and recognition-related tasks and it was theorised that this may be due to a loss of facial expression when they are inverted (Köhler, 1940). Yin (1969) sought to assess whether any given category of mono-oriented stimuli would be subject to this effect, noting that there is some evidence that alphabetic letters are better recognised than their mirror images (Henle, 1942) and that children are sensitive to orientation in the recognition of realistic figures (Ghent, 1960). Yin tested performance in a forced-choice task consisting of a study phase in which they were shown a series of images and a test phase in which a pair of images were shown, one old (from the original series) and one new (previously unseen). In experiment 1 Yin used upright or inverted images of faces, houses, planes, and men in motion. It was found that when presented upright faces were recognised better than other classes of upright stimuli, but when inverted faces were worse than other classes of inverted stimuli. This represents the first demonstration that the inversion effect for faces is greater than that for other type of stimuli (the face inversion effect). In experiment 2 participants were again shown upright or inverted stimuli in the study phase but here they were asked to mentally invert the stimuli before the test phase. In the test phase the stimuli shown were the same images as those in the study phase but presented in the opposite orientation. In this instance faces were the only class of stimuli which exhibited a significant inversion effect. There are two interpretations offered for this, the first being that there is some face-specific mechanism which is particularly sensitive to inversion, and the second being that the easier recognition is when a stimulus is upright, the more greatly it will be impacted by inversion (meaning that the inversion effect for faces may be larger due to better recognition performance when they are upright compared to the other stimuli used in this experiment). To differentiate between these interpretations, experiment 3 replaced the photographic face images with line drawings to eliminate the light and shadow cues, matching the style of the other stimuli which in this case were costumed figures. It was revealed that upright performance for the costumed figures was higher than that for faces but that only faces showed a significant inversion effect, refuting the hypothesis that easier stimuli are more disadvantaged by inversion. Based on this Yin concluded that the data supported the interpretation that faces are special in their processing, making them more difficult to recognise when inverted. Although it is not specified what this face-specific mechanism may be, it was noted that for most classes of stimuli participants reported trying to pick specific features of the image to remember but for faces they attempted to remember a general impression of the image, which they struggled to do when it was inverted. This specificity account was further evidenced by comparisons made between the inversion effect for faces and houses (Valentine & Bruce, 1986) and faces and chairs (Yovel & Kanwisher, 2005); these demonstrated that not only is the behavioural inversion effect greater for faces but also that it is closely associated with the fMRI response in the fusiform face area (FFA) which offers a potential neural mechanism for the face inversion effect.

1.2.2 How acquired prosopagnosia supports the specificity account (Farah et al., 1995; Busigny & Rossion, 2010)

Experimentation on a patient with acquired prosopagnosia (a specific impairment to face recognition as a result of brain damage) has offered further evidence that there may be a face-specific neural mechanism driving the inversion effect. Farah et al. (1995) conducted a case study on a man (LH) with prosopagnosia (with bilateral occipitotemporal lesions and right frontal and anterior temporal lesions) using a matching task, this was chosen over a recognition task due to the profound nature of LH’s impairment and the need to reduce the difficulty of the experimental task. The matching task consisted of 30 pairs of sketched faces with 15 of the pairs being the same and 15 different, each was shown in both the upright and inverted orientation and the task was first tested on healthy participants, and it was shown that the usual inversion effect was found with faster, more accurate performance found for upright compared to inverted faces. When in the subsequent experiments LH was tested however the opposite result was found with performance for inverted faces significantly better than upright. This finding was replicated in additional experiments, once with similar mean scores as the first but a non-significant difference thought to result from a reduced number of trials and again in an experiment with the number or trials increased again which resulted in a significant difference, with inverted faces better than upright. The explanation offered for this by Farah and colleagues is that there is a face-specific mechanism engaged for upright faces which is damaged in LH but still used despite no longer offering an advantage for these stimuli.

There is a key limitation of these experiments in that performance for faces is reported in isolation without comparison to other mono-oriented stimuli. De Gelder and Rouw (2000) found that when tested on non-face stimuli LH showed the same pattern of results with inverted stimuli recognised better than upright indicating that rather than evidencing a face-specific mechanism, the findings of Farah et al. (1995) may reflect an overall visual impairment (a comprehensive overview of LH’s reduced visual capabilities can be found in Levine & Calvanio (1989)). Busigny and Rossion (2010) aimed to address some of the key issues in the previous prosopagnosia literature e.g. discrepancies in the reporting of reaction time data, the type of recognition tasks used, lack of comparison between faces and other objects, and confounding visual impairments. They tested a participant with acquired prosopagnosia named PS who exhibited no deficits in object recognition or other visual abilities outside of the face recognition impairments. All comparisons were made against age-matched healthy controls and reaction time was recorded along with correct responses to allow analysis of any speed-accuracy trade off (this is particularly important in the case of prosopagnosics who may develop analytic strategies to face recognition which take longer than typical face recognition and allow them to achieve closer to normal performance). In experiment 1 they used the Benton face recognition test (BFRT) for which controls showed the expected large inversion effect. PS's performance was reduced compared to controls for upright faces but not for inverted faces in terms of accuracy, and there was no significant inversion effect in this case. This finding was maintained in experiment 2 using a delayed matching task with no difference in either accuracy or reaction time between upright and inverted faces for PS. Comparison against non-face stimuli (cars) in experiments 3 and 4 showed that again for faces controls showed an inversion effect that was not present for PS while for cars PS and controls showed a similar pattern of results with no accuracy inversion effect but a significant rection time inversion effect. This indicates that the deficits to face recognition experienced by PS do not extend to other types of stimuli and may therefore be related to damage to a face specific brain region. An exploration of how PS recognises highly familiar faces was made possible by her work as a kindergarten teacher which provided a set of around 25 faces that she sees very regularly and comparison was made against colleagues with no neurological impairments. The faces were cropped to an oval to remove cues such as hair colour and accessories (which PS reports using to distinguish between individuals while teaching) the familiarity task involved determining whether a given picture was of a child from the kindergarten group. The colleagues performed very highly for the upright faces but much lower for the inverted faces resulting in a large inversion effect. PS however obtained much lower accuracy scores for the upright and inverted faces with no significant difference between the two. Findings across these experiments show than in contrast to Farah et al. (1995) there was no difference in performance between upright and inverted faces, and in no case were the inverted stimuli recognised better than upright. The authors interpret their findings in terms of acquired prosopagnosia disrupting the specific ability to process upright faces as a gestalt (which for healthy controls is disrupted by inversion), thus resulting in an elimination of the inversion effect.

1.2.3 The role of familiarity in the specificity account (Scapinello & Yarmey, 1970)

Yin's (1969) specificity account has also been extended to include additional factors which may influence recognition performance. Scapinello and Yarmey (1970) examined whether familiarity plays a role in the inversion effect for different classes of stimuli (human faces, canine faces, and French architecture) in both immediate and delayed recognition tasks. Their experiment involved an inspection phased in which a set of flashcards equally split between each stimulus group were shown to participants in the upright orientation, half of these were shown once (low-familiarity condition) and half were shown for seven consecutive trials (high familiarity condition). In the subsequent test phase stimuli were shown either upright on inverted (depending on the condition) and half of the flashcards were replaced with new images (still equally split between the three different stimulus classes), participants were then asked to identify which flashcards had been shown in the inspection phase and which were new. The results indicated that for all stimulus types, performance was better in the familiar condition compared to the unfamiliar condition and was reduced in the inverted orientation compared to the upright orientation. Additionally, it was found that regardless of familiarity, human faces were the most greatly affected by inversion, supporting Yin’s (1969) interpretation that the face inversion effect is driven by a face-specific processing/recognition mechanism.

1.3.1 The alternative to specificity: The expertise account

The evidence discussed thus far stands in support of a specificity account of face recognition. There are however clear limitations in terms of the stimuli used for comparison with faces. Humans have such a vast wealth of experience in recognising faces, it is a skill we use every day for the majority of our lives whereas differentiating non-face stimuli such as houses, chairs, architecture, and dogs from one another is not something that most people practice with any regularity. The following research therefore address this disparity and in doing so moves beyond the specificity account in favour of an expertise account. Diamond and Carey’s (1986) work focused on the theory that the face inversion effect is driven by expertise with prototype-defined categories and that it is the type of information contained within the features of the face which poses a problem to a general recognition system when inverted. Distinction is made between 3 types of information: featural information referring to the individual features of the face (e.g., eyes, nose, mouth), first-order configural information referring to the average spatial relationship between the features (e.g., two eyes above a nose above a mouth) and second-order configural information referring to the variations in first-order information in a given exemplar compared to the prototype face. They posited that configural information was likely to be important given that inversion had previously been shown to disrupt the use of configural aspects of faces in a matching task (Sergent, 1984). Diamond and Carey therefore aimed to test hypothesis that the greater inversion effect for faces was due to the use of configural features which are not used in many of the stimuli they had previously been compared to. They can be used for stimuli such as dog faces, stick figures and airplanes but participants do not have expertise in distinguishing these stimuli classes like they do for faces. In their pilot experiment they used the same face stimuli as Yin (1969) and compared them with landscape images in a forced-choice recognition task; the old and new landscape images were paired based on sharing the same isolated features (e.g., mountains, trees, bodies of water) but in different spatial configurations. For both faces and landscapes they found that inversion resulted in significantly reduced performance compared to upright presentation, but that the inversion effect was significantly greater for faces than landscapes. Faces all share a configuration in a way that landscapes do not i.e., two eyes above a nose above a mouth, so for faces first-order configural information is constrained and recognition relies on second-order configural information. Thus, in experiment 2 comparison was made between faces and another stimulus group with first-order configural properties constrained; they relooked at dogs and hypothesised that previous studies had not found a comparable inversion effect to faces because participants are far more expert in using second-order configural information to distinguish human faces. To address this, they used dog experts to assess whether their recognition memory was as disadvantaged by stimulus inversion as human faces are. All participants recognised upright faces better than inverted but only experts recognised upright dogs better than inverted, resulting in an inversion effect comparable to that for faces in the dog expert group. Based on these findings it was concluded that faces are not special due to their features being uniquely represented in memory and therefore sensitive to inversion, but rather suggested that the increased inversion effect is the results of 3 conditions (1: stimuli sharing a configuration, 2: it being possible to distinguish an individual stimulus based on second-order configural features, 3: participants having the expertise to use these features).

1.3.2 Support for expertise: acquired expertise for artificial stimuli (Gauthier & Tarr, 1997)

This expertise account is strongly supported by evidence showing that acquired expertise with an artificial set of prototype-defined stimuli results in a large inversion effect. Gauthier and Tarr (1997) investigated whether sensitivity to changes in configural information is face-specific or can be observed for other stimuli with which have expertise. They noted that faces have a number of characteristics which separate them from previously used stimuli; they have similar features organised in a similar configuration, not just for the exemplars in the experiment but for all known exemplars, faces are also recognised at an exemplar-specific level, and people are highly expert in recognising faces. To address these issues, an artificial class of stimuli named greebles was created to match the configural constraints of faces; they were similarly prototype defined with each exemplar belonging to one of five ‘families’, and one of two ‘genders’. Each family, gender and individual greeble generated was given a nonsense word name which allowed them to test recognition at an exemplar-specific level. Expertise was controlled by training some of the participants on a subset of greebles prior to testing. The experimental design involved the study of a whole individual and then a forced-choice recognition task for the features of the greeble, these were tested in isolation, in the studied configuration and in a transformed configuration (with one of the other features altered to ensure that the transformation was independent of the information required to recognise the target feature) this was conducted on upright and inverted greebles separately for both experts and novices. Based on this design it can be inferred that if the greeble parts are encoded individually, there should be little difference between these recognition conditions but if they are encoded in their configuration (as predicted by the expertise account) then performance should be highest in the studied configuration condition and poorer in the isolated and transformed configuration conditions. Likewise, performance for upright greebles should be more sensitive to configural transformation than inverted greebles, especially for experts. In line with expertise account, experts showed faster, more accurate recognition of the upright stimuli than novices and were also more sensitive to configural transformations, with faster recognition only in the studied configuration condition. That these findings can be shown in non-face stimuli indicates that configural sensitivity is not face-specific but rather the result of a general mechanism that is honed by experience with homogeneous stimuli and manifests as an advantage when the stimuli is presented in its typical form (e.g., upright faces) and a more prominent disadvantage when presented in a different configuration (e.g., inverted faces).

1.4.1 Perceptual learning and the MKM model (McLaren, 1997, McLaren, Kaye & Mackintosh, 1989)

McLaren (1997) offers an explanation of why specifically disruption to configural information may produce the differential effects that it does for upright and inverted stimuli. The underlying theory relates to latent inhibition and perceptual learning and posits that expertise with a prototype-defined stimulus reduces the relative salience of elements common across all exemplars and increases the relative salience of elements unique to specific exemplars (which are those that best enable discrimination between exemplars). This theory is derived from the associative mechanism outlined in the MKM model by McLaren, Kaye, and Mackintosh (1989). This model explains perceptual learning in terms of the differential latent inhibition of the common elements representing the stimuli. Within the model stimuli are represented as ‘units’ corresponding to sets of features which exist in an error-correcting associative network in which the error term delta (Δ) can modulate the salience of the units depending on their level of activation and in doing so controls learning. During a given stimulus presentation a subset of units will be activated, which ones may depend on the conditions of the presentation such as angle, lighting, and the presence or absence of other stimuli. The subset of units simultaneously activated become associated with each other within the network as learning occurs about the stimulus. When a stimulus is presented repeatedly (as in the case when expertise is high) the elements that are found in nearly every exemplar of the stimulus will be consistently associated with one another. The error term controls learning by modulating the salience of a given unit based on how well it is predicted by other active units in the model; when it is well predicted the error term remains low and thus its salience is also low (latent inhibition), when it is not well predicted the error term is high and thus its salience is also high and it will readily form new associations (learning). Therefore, when a prototype-defined stimulus is presented many times the salience of the features that are common across exemplars and are frequently activated together and thus represented by very well predicted units will be low. However, features that are unique to an individual exemplar are activated fairly infrequently across different stimulus presentations, so the salience of these unique features is high. In this way we can see how expertise might lead to greater reliance on second-order configural information (unique between exemplars) and reduced salience of first-order configural information (common between exemplars) and enhances performance for familiar upright stimuli. In terms of the face inversion effect this can explain why we are so superior in our recognition of upright faces (with which we are experts) compared to inverted faces (with which we have not gained this perceptual expertise). The predictions made by the MKM model in terms of perceptual learning are empirically supported by McLaren’s (1997) work involving prototype-defined categories of checkerboards. Similar to Gauthier and Tarr's (1997) greebles these are artificial stimuli for which expertise can be controlled, they are not commonly experienced in daily life and therefore there is no reason to expect participants to have a base level of expertise with them. They are also non-mono-oriented meaning that rather than having set upright and inverted orientations, this is defined by the orientation they are initially presented to participants in during training. In experiment 1 two prototype checkerboards were generated for each participant consisting of 16x16 black and white squares, these then formed the basis of two distinct categories of checkerboards. Each stimulus in a given set was generated by varying rows from the original prototype such that on average they shared around 80% of squares with the prototype but could have considerably varying configurations from one another. Participants were trained on two prototype categories during a pre-exposure phase through trial-and-error categorisation of the checkerboards and one of these now familiar categories was then used alongside a novel prototype category in a discrimination phase with half upright and half inverted stimuli (relative to the orientation shown during pre-exposure). It was demonstrated that it was possible to produce an inversion effect for these stimuli only when they were drawn from familiar (pre-exposed) prototype-defined categories. Experiment 2 extended this from a discrimination paradigm to a recognition paradigm using a matching task and again it was found that for familiar prototype-defined checkerboards there was a significant inversion effect. These results are in line with the MKM model’s prediction that expertise with the upright stimuli enhances the discrimination between exemplars and that this advantage is lost when they are inverted.

1.4.2 Checkerboards as an analogue to the face inversion effect (Civile, Zhao et al., 2014)

The checkerboard stimuli used in McLaren (1997) were subsequently used by Civile, Zhao, et al., (2014), who investigated them in the context of an old/new recognition task to provide results comparable to those found for the face inversion effect which typically uses this task. The aim being to demonstrate that recognition for faces shares a mechanism of feature salience with non-face stimuli for which we have expertise. Experiment 1a began with the categorisation used in McLaren (1997) in which participants sorted checkerboards from two different prototype categories through trial-and-error. In this case four checkerboard prototypes were created, consisting of 16x16 black and white squares, and sharing 50% of these squares with each of the other prototypes. The creation of individual exemplars was achieved by randomly altering 48 of the squares (with around half of these changing from black to white or vice versa). Following categorisation of checkerboards derived from two of these prototypes, participants engaged in a study phase in which they attempted to memorise checkerboards from one of the categories used in the categorisation task (although not the same exemplars that were used here) and one novel category. Half of these were presented upright (i.e., in the same orientation as in the categorisation task) and half were inverted. The recognition phase then involved the old exemplars from the study phase (in the orientation studied) and new exemplars not seen in the study phase but from the same two prototype-defined categories used in that phase (half upright, half inverted), and participants were asked to respond to each indicating whether or not they had seen it in the study phase. A robust inversion effect for checkerboards drawn from the familiar category was found while none was found for those in the novel category. While this effect was significant the difference in performance between exemplars taken from the familiar and novel categories, specifically for upright stimuli, was not in this case. In order to demonstrate that this difference in inversion effect results from familiarity with a prototype-defined category, experiment 1b altered the checkerboard stimuli such that they were no longer derived from a prototype. The experimental design was the same but in this case the checkerboards were “shuffled’ with two rows swapped and then an additional row swapped with one of these to maintain the average of 24 squares changing as in experiment 1a. In this version of the experiment no inversion effect was found for either familiar or novel checkerboards indicating that the inversion effect is dependent prototype-defined stimuli. Experiment 2 aimed to make recognition of the checkerboards easier in order to increase produce a stronger inversion effect for the familiar exemplars. This was achieved by clumping the black and white squares together more by making the probability of the colour for a given square dependent on the colour of the squares adjacent to it. Having more black squares nearby increased the likelihood of a square being black and the same for white squares. In this way they were able to maintain the 50% black and white squares and prototypes sharing 50% of the same squares with one another whilst also making recognition easier. This resulted in an increased inversion effect for the familiar checkerboards, confirming the results of experiment 1a. In experiment 3a these same clumpy checkerboards were used while in experiment 3b they were shuffled in the same manner as the checkerboards in experiment 1b. This served to replicate the comparison between prototype-defined and non-prototype-defined checkerboards with the easier to recognise stimuli in a single experiment. Here again the was a significant inversion effect for familiar prototype-defined checkerboards but not for novel or non-prototype-defined checkerboards, demonstrating that both of these factors are required for the inversion effect to occur. Taken together this series of studies supports expertise playing a role in the inversion effect and also lends credence to the MKM model (McLaren, Kaye, et al., 1989); success in the categorisation phase relies on the common features of the stimuli being identified (as this is indicative of group membership), and thus these common features become well associated with one another and lose salience as perceptual learning occurs. The features unique to each exemplar however are still forming associations and thus remain highly salient.

1.5 How different types of information impact the inversion effect

The literature surrounding how different types of information influence the inversion is not entirely clear, with different research offering conflicting evidence about the specific types of information that might be causal. In the following section I will explore some of this literature looking into types of information both individually and in combination. There is an additional debate in this body of work surrounding the nature of holistic information, how it is defined and the role it plays which will also be explored.

1.5.1.1 Featural and configural manipulations (Leder & Bruce, 1998)

The role that configural information plays in the inversion effect has been assessed through a variety of stimulus manipulations, Leder and Bruce (1998) independently manipulated featural and configural information of faces to increase distinctiveness and directly compare the how these types of information are impacted by inversion. If the face inversion effect is due to a face-specific mechanism which is generally impaired by inversion, then there should be no difference between these types of information but if it results from specialised impairment to configural information then it would be expected that featural and configural would differentially affect both the inversion effect and apparent distinctiveness. The stimuli in this study began with faces which has previously been rated as average in distinctiveness, which were then altered to increase distinctiveness; the featural manipulation (D-local) involved thickening and darkening the eyebrows and the mouth, broadening the nose etc., while the configural manipulation (D-rel) involved changing the size of the gaps between the two eyes and the nose and mouth. Participants rated distinctiveness of all 3 versions of each face (original, D-local, D-rel), half rating them in the upright orientation and half in the inverted orientation. The overall findings from the studies revealed that, as expected, when presented upright D-rel and D-local faces were rated as more distinctive than their original counterpart. However, when inverted only the D-local show this effect while it is lost entirely for the D-rel versions. A similar pattern of results was shown when these stimuli were tested for recognition performance with D-local and D-rel recognised better than originals in the upright orientation, and a significant reduced in the effect for D-rel stimuli when inverted. These results support the idea that inversion disrupts the processing of specific sources of information differentially, with configural information more sensitive to inversion than featural information.

1.5.1.2 Configural information in the inversion effect for dot patterns (Tanaka and Farah, 1991)

The role of configural information in the inversion effect was investigated by Tanaka and Farah (1991). They used as stimuli dot patterns, some of which were derived from a prototype and as such shared a spatial configuration with one another and some which were not. Given that the isolated features of a dot pattern are all the same (i.e., an individual dot), discrimination between exemplars is only possible through the spatial configurations. If we apply Diamond and Carey's (1986) theory, dot patterns derived from the same prototype have in common their first-order configuration and thus discrimination between them relies on second-order configural information, whereas those not defined by a prototype are discriminable based only on first-order configural information. Participants were trained on either first-order patterns or second-order patterns which had been assigned male and females names for the purposes of identification. In this learning phase 6 patterns were presented, and participants were asked to identify them by name, the phase continued until each pattern had been correctly identified twice without error. In the test phase these 6 dot patterns were presented in either the upright or the inverted orientation and participants were again asked to identify them by their assigned names. If second-order configural information plays a causal role in the inversion effect, then we would expect to see that recognition for second-order patterns was more greatly impaired by inversion than first-order patterns. However, this was not supported by the results of this study, instead both sets of patterns displayed an inversion effect, and the magnitude of this effect did not differ between groups. This result persisted in experiment 2 for which the second-order patterns were designed to share an even greater proportion of their first-order spatial configurations (increasing participants reliance on second-order information). On the basis of this evidence, it was concluded that second-order configural information is not inherently more sensitive to inversion that first-order configural information, refuting the explanation of the face inversion effect offered in Diamond and Carey (1986). There are however some critical aspects of this study to note when considering this interpretation; in contrast to other experiments investigating the role of second-order configural information in the inversion effect (e.g. McLaren, 1997) the configurations in this case are not pre-exposed and there is therefore no opportunity for participants to develop expertise with the second-order configural information which is a crucial component of both Diamond and Carey’s (1986) explanation and the MKM model of perceptual learning offered by McLaren, Kaye and Mackintosh (1989). A further issue in terms of the MKM model is that it relies on the common elements between exemplars being well predicted by one another, thus increasing the relative salience of the unique elements, however for both the first and second-order dot patterns the position of each dot is changed in each exemplar and there are therefore very few common elements becoming associated with and predicted by one another and the relative salience is not modulated. With these factors in mind we can explain some of the discrepancy between these findings and the literature which has found that second-order configural information contributes to the inversion effect.

1.5.2.1 Featural information and the inversion effect (McKone & Yovel, 2009; Rakover & Teucher, 1997)

In addition to configural information, the role of featural information in the inversion effect has also been investigated and this offers some alternative findings to those previously presented which suggest the configural information drives the inversion effect. McKone and Yovel (2009) conducted a meta-analysis investigating whether disruption to featural information can impact the face inversion effect in a similar manner to configural disruptions. They found that altering the shape of individual features impacts the inversion effect to an equivalent extent to configural manipulations. Evidence that featural information alone can induce an inversion effect can be seen in Rakover and Teucher (1997). They used a Photofit kit to divide faces into five distinct dimensions (hair and forehead, eyes and eyebrows, nose, mouth, and chin). For each of these dimensions they selected 20 different versions from the Photofit kit which varied in size, colour, shape etc. Their experimental design consisted of a study phase in which 10 versions of a given dimension were presented in random order, half of participants saw these in the upright orientation and half inverted. In the follow test phase, the original 10 were shown intermixed with the previously unseen 10 versions, here depending on their condition participants may have seen the features in a congruent or incongruent orientation to that shown in the study phase, resulting in 4 conditions (Upright-Upright, Upright-Inverted, Inverted-Upright, and Inverted-Inverted). Each different dimension was shown to each participant in successive experiments. An overall inversion effect was found with performance in the U-U group higher than any other condition. This indicates that contrary to Diamond and Carey’s (1986) explanation, an inversion effect is not solely dependent on configural information but can also be obtained with only the isolated features of the face.

1.5.2.2 Is configural information necessary for the face inversion effect?

Scrambled faces (Civile et al., 2014)

Civile et al. (2014) aimed to establish whether first and second-order configural information are required to produce an inversion effect in whole faces. They compared normal faces and scrambled faces (for which first and second-order information was completely disrupted), asserting that if the inversion effect is eliminated for scrambled faces, this would support Diamond and Carey's (1986) position, but a remaining inversion effect would indicate that configural information does not play a sole role in the inversion effect. In experiment 1a the scrambled faces were created by repositioning the features of the face (the eyes and eyebrows, the nose, the mouth, and the ears), in order to ensure complete disruption to the typical first-order configuration each eye/brow and ear is treated as an individual feature. One of these features is first moved to the forehead (chosen because it provides the largest free space on the face), a second feature is the moved to take the position of the first and so on until all features have been repositioned. The sequence of features moved was consistent across exemplars, meaning all scrambled faces shared a prototypical configuration, although not the one typically seen for faces. An old/new recognition task was employed again here, in the study phase normal and scrambled faces were intermixed with half of each face type presented in the upright orientation and half in the inverted orientation. These faces were then shown again in the recognition phase (in the same orientation as they were originally presented) along with a series previously unseen faces (also equally split between normal, scrambled, upright and inverted), and participants indicated whether or not each face had been seen in the study phase. The results revealed that while overall performance for scrambled faces was reduced compared to normal faces, disruption to both types of configural information did not differentially impair performance for upright and inverted exemplars and as such did not result in a significant reduction of the inversion effect. This result was maintained in experiment 1b which smoothed the normal faces to more closely match the light/shadow cues and local feature information in the scrambled faces and tested different scrambling configurations to ensure that the result was not due to the difficulty of the specific scrambled configuration in experiment 1a. These findings therefore are not consistent with the explanation of the face inversion effect offered by Diamond and Carey (1986). In experiment 2, Civile and colleagues extended their scrambling manipulation to investigate whether the individual features of the face may play a role in the inversion effect. They created stimuli known as “50% feature-inverted and scrambled faces” using the 4 categories of scrambled faces created in 1b and rotating half of the features (one eye/brow, one ear and either the nose or mouth) by 180° such that no matter the overall orientation of the face, half of the feature are always inverted. They employed the same old/new recognition task and found that with all configural information entirely disrupted *and* the single feature orientation manipulated the inversion effect no longer reached significance and was thus entirely eliminated. These findings support that of Rakover and Teucher (1997) and support featural information as a primary component of the face inversion effect.

1.5.2.3 Single feature orientation new Thatcherised faces (Civile et al., 2016)

In light of the previous findings, to investigate whether second-order configural information is at all necessary for the production of the inversion effect it is necessary to manipulate this information while single feature orientation is controlled for. Civile et al. (2016) based their manipulation on the ‘Thatcher illusion’

(Thompson, 1980) with some key alterations. The original Thatcher illusion was created using a poster of British PM Margaret Thatcher and rotating the eyes and mouth 180°, when this image is inverted these changes are very difficult to detect, even when presented alongside the unaltered poster. However, when the posters are presented in the upright orientation the manipulation is immediately obvious and the result is that the ‘Thatcherised’ image looks extremely bizarre. This effect has been utilised to investigate the inversion effect previously and is useful in that the featural and second-order configural information or disrupted while first-order configural information is relatively unaltered. There is however a potential limitation involved in rotating both of the eyes due to their relatively high salience in face recognition (Ellis et al., 1979; Haig, 1984). Civile et al. (2016) therefore revised the Thatcher illusion in a similar manner to the 50% feature-inverted faces in Civile et al. (2014) but this time beginning with a normal face. As before half of the features (one eye/brow, one ear and either the nose or mouth) were rotated 180°. This controlled single feature orientation as half were always upright/inverted no matter the overall orientation of the face, second-order configural information was somewhat disrupted as the rotation results in the variations in the spatial configuration being altered slightly. First-order configuration on the other hand remains relatively unaltered by the manipulation. In experiment 1 comparison between normal and new Thatcherised faces in an old/new recognition task revealed that the manipulation resulted in a significant reduction in the inversion effect although it was not eliminated entirely. Civile and colleagues note some tension between these findings and that of Civile et al. (2014) in which there is no significant reduction in the inversion effect for scrambled faces, which visually at least have much greater disruption to second-order configural information, and they also showed complete reduction of the inversion effect when 50% of the features are inverted but some of this remains in the current study. Follow-up experiments focused on a direct comparison between the scrambled faces and 50% feature inverted scrambled faces used in Civile et al. (2014) and the new Thatcherised faces. The results here indicated confirmation that scrambled faces produce a robust inversion effect, the new Thatcherised faces showed a reduced but still significant inversion effect, and the 50% feature-inverted scrambled faces showed no significant inversion effect. Taken together the findings demonstrate that when feature orientation is controlled for, leaving first-order and some second order configural information intact can produce an inversion effect, thereby supporting the assertion that both featural and first-order (and possibly second-order) configural information are causal to the inversion effect.

1.6.1 A holistic account: parts and wholes (Tanaka & Farah, 1993)

One possible explanation for the somewhat conflicting evidence regarding how different information influences the inversion effect and the above finding that a combination of disruptions is required to eliminate the inversion effect is that featural and configural information interact to form a holistic representation. This is the basis of the theory set forth by Tanaka and Farah (1993) that faces are represented more holistically than other classes of stimuli. To investigate this hypothesis, they designed a series of experiments in which a part of a given stimulus is recognised in the context on the whole stimulus or in isolation. The prediction here being that recognition for a part of a face will be disproportionately easier when shown in the whole face compared to in isolation, relative to recognition of the parts and wholes of other classes of stimuli. Across three experiments normal faces were compared to scrambled faces, inverted faces, and houses; in each case a learning phase took place in which participants attempted to memorise the name-stimulus association (for houses the name corresponded the owner of the house). A forced-choice recognition task followed in which pairs of stimuli were presented consisting of either the whole stimulus or an isolated part and participants were asked to identify which image corresponded to the target name. The results in each experiment supported the predictions made, with isolated parts of faces recognised worse than whole faces. This result did not extend to any other class of stimuli used with no disadvantage seen for parts over wholes, substantiating Tanaka and Farah's hypothesis that faces are represented more holistically than other stimuli.

1.6.2 Features and their configuration (Tanaka & Sengco, 1997)

Additional support for this theory is provided by Tanaka and Sengco (1997); they argued that if faces are represented holistically through a combination of featural and configural information then changes in configural information should affect the recognition of the individual facial features. Their stimulus manipulation involved the creation of 2 different configurations by either increasing or decreasing the inter-ocular distance of the faces. Participants studied faces in one of these configurations and then recognised parts of the faces either presented within this same configuration, within a new configuration (the one not presented in the study phase), or in isolation. The holistic account posits that participants would be expected to recognise features better in the original configuration, where the second-order configural information remains unaltered from the study phase, compared to the new configuration where it was altered. Recognition for the features of the face not involved in the configural manipulation (the nose and mouth) would also be expected to be impaired when seen outside of the original configuration if faces are represented holistically. In experiment 1 it was found that, in line with the holistic account, features were recognised significantly better in their old configuration compared to the new configuration or in isolation. It was also found that the nose and mouth were also subject to this effect despite being unaltered by the configural manipulation. In experiment 2 this paradigm was extended to include inverted faces and comparison between these, and the normal faces showed that the configural manipulation had no effect on the recognition of features in the inverted orientation, indicating that manipulating second-order configural information affects only the recognition of facial features in upright faces. Taken in summary these studies show clear support for a holistic theory of facial representation.

1.6.3 Holistic processing of photographic negatives (Hole, George & Dunsmore, 1999)

Hole et al., (1999) theorise that holistic processing may in essence be a type of configural processing which binds facial features into a ‘gestalt’ rather than acting in isolation. This argument is based on their work with the chimeric face effect, this relates to a category of face stimuli for which the top half of one face is paired with bottom half of another; when upright, recognition of the top half of the face is made more difficult by the configural information between the two halves, when inverted however, this configural information is disrupted making the bottom half easier to disregard in favour of better recognition for the top half. They used photographic negatives of faces to study this effect, similar to inversion, negative contrast in faces is also shown to be impair recognition performance compared to faces with normal contrast (positives). The process by which it does this however is argued to be different than that for inversion given that the configural relationship of features should not be altered by this transformation, making it seem likely that these manipulations are affecting different elements of the process of face recognition. In experiment 1 Hole et al. presented participants with pairs of chimeric faces and asked them to identify whether the top halves of those faces were identical (the bottoms were always different). Stimuli in this task was split into four categories, upright positives, inverted positives, upright negatives, and inverted negatives. Each participant saw each set of stimuli in a blocked design. Positive faces should produce the standard chimeric face effect with better identification of identical pairs in the inverted condition compared to the upright condition as a result of reduced interference from configural processing. If negatives are also subject to configural processing this effect should also be found in those conditions. The results from experiment 1 are in line with these expectations, negatives appear to produce the chimeric face effect in much the same was as positives do with reaction time data and (to a lesser extent) accuracy data showing that inverted chimeras benefit from disruption to configural processing. This provides initial evidence that photographic negative faces maintain the configural aspects of face recognition seen in normal faces and that the negative filter is therefore targeting a different aspect when reducing recognition performance. Experiment 2 addressed the criticism that pairs of identical pictures may be identified on the basis of some pattern matching mechanism of the features which would not translate to real life recognition of faces. In this instance participants were shown pairs of chimeras in which one was a full-face view and the other was a ¾ profile view, this meant that identification must be made on the basis of the individual in the photograph being the same rather than simply the content of the images. The findings here further support the idea that negatives of faces produce the chimera face effect and thus are subject to some form of configural processing. Some of the previous evidence relating to negatives has identified difficulty in detecting configural manipulations (e.g., Thatcherisation) with these stimuli compared to positive faces (Lewis & Johnston, 1997). Additionally, the chimeric face effect and the inversion effect have been shown to have a different developmental time-course with the chimeric face effect comparable in young children and adults while the inversion effect grows more pronounced with age (Carey & Diamond, 1994). These differences seem to suggest that there is more at play here than both effects reflecting the same process of configural disruption. In their paper Cary and Diamond (1994) offer the explanation that the chimeric face effect relates to a form of holistic processing that operates purely to establish that the perceived stimulus is a face, while the inversion effect relates to a more specialised form of configural processing which develops with experience. On the basis of this, Hole et al. (1999) theorised that configural processing (of the kind described when discussing first and second-order information) relates to specific facial details and may be used in the recognition of individual faces while holistic processing can be elicited by any stimulus that corresponds to the basic plan of a face.

1.7 Neurostimulation: The modulation of the face inversion effect through transcranial direct current stimulation

Having established some of the evidence relating to the mechanisms responsible for producing the face inversion effect, I now move on to discussing how it can be modulated through the application of non-invasive neurostimulation using transcranial direct current stimulation (tDCS). tDCS typically involves two electrodes of different polarities (an anode and a cathode) placed on the scalp through which small amplitudes of electricity are passed. This serves to induce changes in the electrical activity of the neurons which in turn can modulate the response threshold of the neurons stimulated by tDCS. Historically, this has resulted in anodal stimulation improving performance and cathodal impairing performance (Nitsche & Paulus, 2000). However, the following studies using a specific tDCS montage developed initially by Ambrus et al. (2011) demonstrate that this effect is not ubiquitous and that the effects of tDCS can depend on the task being undertaken and the area stimulated.

1.7.1 The tDCS set-up and its effect on a categorisation task (Ambrus et al., 2011)

Ambrus et al. (2011) used a variation of a prototype distortion task with the aim of operationalising the acquisition of category learning as distinct from performance in category recognition. tDCS was delivered in order to establish specific regions as playing a causal role in this process, the montages used involved anodal or cathodal stimulation over the left DLPFC, with the return (opposite electrode) at Cz, this was applied at 1.0mA, for 10 minutes. The left DLPFC was chosen for this experiment based on fMRI data from Seger et al. (2000) which was one of the few studies that had focused on the acquisition phase of categorisation tasks rather than the testing phase. This data indicated that while the right DLPFC was active in the early stages of a category learning task for both high and low performers, the left DLPFC is active only in later stages of learning and only for those who performed highly in the task. This provided evidence in support of the notion that the right region may be involved in stimulus processing and visual reasoning whereas the pattern of activation in the left area is indicative of learning related change and supports evidence that the left

DLPFC plays a role in gaining expertise in categorisation. Control groups received sham stimulation, which was delivered at the same level, but ramped down after 30 seconds rather than continuing for the full 10-minute duration (sham stimulation of this kind is widely accepted to have no effect on performance). Stimulation was delivered for 8 minutes prior to the beginning of a training phase for the prototype distortion task. The stimuli used were dot patterns, with a prototype generated which was then used as the basis for high or low distortion dot patterns, depending on the level of variation in dot placement compared to the prototype. During training participants saw high and low distorted patterns intermixed and then at test high and low distortions were presented alongside patterns not derived from the prototype and were asked for each pattern whether it belonged to the category presented during training. In line with Nitsche and Paulus (2000) it was predicted that anodal stimulation would improve performance while cathodal stimulation would reduce performance. However, their results demonstrated the opposite effect, anodal stimulation resulted in impaired categorisation accuracy compared to sham while cathodal stimulation did not provide significantly different results in terms of categorisation performance compared to sham. The prototype effect (better categorisation of the prototype compared to distortions, despite not being presented during the training phase) on the other hand, was present in then sham condition but not in any of the active conditions.

1.7.2 tDCS and the prototype distortion effect- an explanation from the MKM feature salience model (McLaren et al., 2016)

McLaren et al.’s (2016) work built on the findings on Ambrus et al. (2011) and offered an interpretation based on the MKM feature salience model of perceptual learning.

They posit that anodal tDCS over Fp3 is able to change the modulation of salience based on prediction error, in the MKM model this modulation of feature salience allows greater learning about poorly predicted (unique) features of a given stimulus. Disrupting the modulation of feature salience results in a reduction in learning about the unique features and an increase in generalisation. Expanding on the previous work to examine whether the data aligns with the MKM model, experiment 1 of McLaren et al. re-examined the tDCS procedure with a checkerboard categorisation task for which three distinct prototype categories were generated (A, B and C). Participants were first asked to separate exemplars into their categories in a trial-and-error task where they were presented one checkerboard at a time in random order, in the subsequent test phase 10 exemplars from each category which had been shown in the training phase, 10 exemplars from each category which has not been shown, and the 3 original prototypes were presented and participants were again asked to categorise these. The tDCS procedure was derived from that of Ambrus et al. (2011) using a similar Fp3 active channel, the current was increased from 1.0mA to 1.5mA to maximise their chances of seeing an effect on categorisation. Participants were split in anodal, cathodal, and sham stimulation conditions; in the anodal and cathodal conditions stimulation began 1.5 minutes before the categorisation task and continued for 10 minutes, while in the sham condition 30 seconds of stimulation was delivered prior to the start of the categorisation task. For each condition mean accuracy scores for the exemplars was subtracted from mean accuracy scores for the prototype. In line with Ambrus et al. (2011) the results showed a significant prototype effect in the sham and cathodal conditions and a significantly reduced prototype effect (such that it was no longer significant) in the anodal condition, with this difference driven by decreased performance for the prototypes in the anodal condition. Experiment 2 continued with this tDCS procedure and instead used the categorisation training seen in McLaren (1997) and Civile, Zhao, et al. (2014) with 2 different prototype categories to allow them to explore the effect that tDCS might have on these perceptual learning experiments. Across experiments 2a and 2b, anodal stimulation was compared to sham control and cathodal control. The results in this case are different; here in the anodal condition there is a strong prototype effect that appears to emerge as a result of reduced performance for the exemplars while the prototypes remain unaffected. Neither control condition however produced a significant prototype effect. Key to McLaren et al. (2016)'s interpretations of this is that the critical difference between the tasks in experiments 1 and 2 is the number of prototype categories. They argue that the prototype effect seen for controls in experiment 1 is due to the increased generalisation between categories (due to the greater number of them) being counteracted by perceptual learning. tDCS increases this generalisation and decreases perceptual learning which negatively impacts the prototype effect. In experiment 2 however, there is less generalisation between different categories to start with which should make categorisation easier and is helped by perceptual learning. Now tDCS again reduces perceptual learning and promotes generalisation but it is the former effect that dominates this time and the prototype effect increases. It is this balance that they argue shifts between the 2 category and 3 category tasks such that for the former generalisation across categories is dominant while for the latter the strengthening of the perceptual learning effect takes over, explaining the results from the control groups. They further theorise that the effect of anodal tDCS is to reduce perceptual learning and thereby enhance generalisation, this therefore results in the reduction of the prototype effect in experiment 1 as the already high generalisation is increased and the perceptual learning that was acting to figureit is reduced.

1.7.3 tDCS on checkerboard categorisation (Civile, Verbruggen et al., 2016)

We return now to the checkerboard stimuli used in McLaren (1997) and Civile, Zhao et al. (2014), which became the focus of tDCS work for Civile, Verbruggen et al. (2016). They developed an adaptation of the procedure used by Ambrus et al. (2011) with the aim of testing whether it can also modulate perceptual learning as defined in their previous experiments. Checkerboards seem an ideal stimulus to begin this exploration with given that like the dot patterns in Ambrus et al. (2011) they are artificially created, derived from prototype-defined categories, and require pre-exposure in a categorisation task. Experiment 1 investigated the effect of tDCS on the old/new recognition task for checkerboards. The checkerboard stimuli were those previously seen in Civile, Zhao et al. (2014). The experimental design from this study was also used with participants first engaging in trial-and-error categorisation phase, followed by the old/new recognition task. There were slight modifications made to the tDCS procedure in this case, stimulation was delivered over the DLPFC at 1.5mA for 10 minutes, with the return this time positioned above the right eyebrow. During anodal stimulation the current ramped up to its peak of 1.5mA for the first 5 seconds and continued for the duration of the categorisation task, at the end of 10 minutes stimulation faded out for the final 5 seconds. This was compared to sham stimulation which again lasted for only 30 seconds and was completed prior to the start of the categorisation task. Results from the old/new recognition task revealed that in the sham condition familiar checkerboards produced the previously demonstrated inversion effect, however, in the anodal condition this effect was eliminated entirely due to impaired performance for the upright exemplars. In terms of perceptual learning this result can be seen as the tDCS reducing or removing the advantage usually conferred to upright familiar stimuli as a result of perceptual learning making discrimination easier. Experiment 2 followed a similar procedure to experiment 1 but in this case rather than using sham as their control, cathodal stimulation was used. The tDCS montage here was exactly the same of for anodal stimulation except that the cathodal electrode this time was placed over the DLPFC and the anodal acted as the return channel on above the eyebrow. This was expected to produce an inversion effect similar to or larger than the one for sham. The results showed that indeed the familiar checkerboards in the cathodal condition produced an inversion effect that was not present for the novel checkerboards, while in the anodal condition there is again a significant reduction in the inversion effect as a result of impaired performance for upright familiar checkerboards. Taken in combination, these findings demonstrated that anodal tDCS over the DLPFC eliminates the inversion effect for checkerboards by selectively reducing performance for familiar upright exemplars and provides evidence for a reversal of perceptual learning.

1.7.4 tDCS on the face inversion effect (Civile, et al., 2018)

The logical follow-on for this body of tDCS work was to test the effect of anodal stimulation on face recognition and establish whether the effects seen in checkerboards translate to faces which would strengthen the evidence that perceptual learning is involved in face recognition. Civile, McLaren and McLaren (2018) employed the same old/new recognition task previously described but with sets of face stimuli. They conducted a series of experiments differing in tDCS procedure but with the same behavioural design. The study phase consisted of 128 faces (half male, half female) presented intermixed for participants to try to memorise. The recognition phase included these 128 old faces along with 128 novel faces and participants were asked to identify which has been seen in the study phase and which had not. A double-blind procedure was used to administer either anodal or sham tDCS in experiment 1, set-up followed the same parameters as Civile, Verbruggen et al. (2016) with the active electrode over the DLPFC at Fp3 and return above the right eyebrow at Fp2. Anodal stimulation was again delivered at 1.5mA for 10 minutes in the active condition with 5-second ramp up and fade out, this took place during the study phase. In the sham condition 30 seconds of stimulation was delivered at 1.5mA followed by 0.1mA pulses (3ms peak) delivered periodically over 10-minute duration. The results showed the typical face inversion effect in the sham condition and a significantly reduced inversion effect in the anodal condition, driven by reduced performance for the upright stimuli. These results were replicated in experiment 2 which used the same tDCS procedure and behavioural design. These combined findings support the evidence provided by Civile, Verbruggen et al. (2016) regarding the tDCS-induced effects for checkerboards and confirms that these effects are analogous between checkerboards and faces. Experiment 3 tested the same behavioural paradigm with a different positioning of electrodes in the tDCS procedure (active control). This served to investigate whether it is the specific Fp3 stimulation that impacts the inversion effect. For this purpose, they selected the right-Inferior Frontal Gyrus (rIFG) as the site for the anodal electrode which had previously been shown to be effective in other tasks (e.g., go/no go tasks, see Cunillera et al., 2014; Jacobson, Javitt, Lavidor, 2011; Stramaccia et al., 2015) but had not been tested on perceptual learning tasks. The return electrode was place over the left eyebrow at Fp1 which kept the distance between electrode roughly the same as in the Fp3-Fp2 paradigm. In all other respects the tDCS procedure was the same as before with anodal stimulation compared to sham. The results showed that this set-up did not produce the same effects as previously seen; there was no reduction seen in the inversion effect and the results did not differ significantly from sham stimulation. This indicates that it is the specifically stimulation over Fp3 that induces disruption to perceptual learning and thus reduces the inversion effect for faces and checkerboards. Across the wider literature these findings have replicated extensively (Civile, et al., 2019; Civile, Waguri et al., 2020; Civile, Cooke et al., 2020; Civile, McLaren et al., 2020; Civile, Quaglia et al., 2021; Civile, McLaren et al., 2021; Civile & McLaren, 2022; Civile, McLaren et al., 2023), providing a substantial body of evidence to offer support for this conclusion.

1.7.5 tDCS on the composite face effect, active control (Civile, McLaren et al. 2021)

Further research has investigated other variations in the tDCS procedure through the use of different active controls. In Civile et al. (2018) active control is achieved by changing the position of both the anodal and cathodal (return) electrodes. In Civile, McLaren et al. (2021) active control was again employed but, in this case, only the anodal electrode was repositioned. The perceptual learning task used in this case was a composite face task; this involves stimuli comprised of the top half of one face and the bottom face of another (composite faces). The composite face effect (also referred to as the chimeric face effect) refers to the phenomena in which these composites are perceived as novel holistic configurations when the two halves are aligned and presented upright; under these conditions two identical top halves are difficult to perceive as such when paired with different bottom halves and vice versa.

However, this effect disappears when the halves are misaligned or when the composites are presented in the inverted orientation. Civile et al. (2021) used a matching task in which a target composite was shown followed immediately by a test composite with participants responding to whether the top halves of the target and test faces was the same or different. Trials were split into four groups based on whether the composites were aligned or misaligned (meaning that the two halves had been shifted with one moving left and the other right), and whether the target and test composites were congruent (either the same top half and bottom half or different top and bottom halves) or incongruent (the same top half paired with a different bottom half or a different top half paired with the same bottom half). Experiment 1a used only upright faces and experiment 1b included upright and inverted (with orientation always the same at target and test) In both experiments the established Fp3-Fp2 tDCS montage was used with both anodal and sham stimulation following the same intensity and durations described in Civile et al. (2018). The results showed that tDCS had no effect on the composite face effect although performance for upright stimuli was impaired by the tDCS procedure, experiment 1b again demonstrated the reduction in the inversion effect caused by anodal stimulation. Experiment 2 replicated and extended experiment 1a with the inclusion of an additional tDCS condition. In addition to the Fp3-Fp2 and sham montages an active control was added with the active stimulation site at PO8, this was chosen based on previous evidence that showed the N170 ERP component and its modulation in response to normal/distorted faces and prototype-defined stimuli to be greatest at this channel. As an active control the intensity and duration of stimulation delivered in the PO8-Fp2 set-up was the same as that for the FP3-Fp2 set-up. The results again revealed that while the tDCS had no effect on the composite face effect, anodal stimulation in the Fp3-Fp3 condition did impair performance for the upright faces compared to sham and in this case compared to the PO8-Fp2 condition (which did not differ significantly from sham). These findings provide support for their perceptual learning account by demonstrating the predicted decline in performance for the upright faces. Importantly in terms of the modulation of perceptual learning they also demonstrate that this this tDCS-induced reduction was only seen in the Fp3-Fp2 condition, indicating that it is specific to this paradigm.

1.7.6 Differences in the effect of tDCS on the inversion effect for faces and checkerboards (Civile, Quaglia et al. (2021)

It has been demonstrated extensively throughout the literature that anodal tDCS is able to reduce the inversion effect for faces and checkerboards, however, there is a notable difference in the tDCS-induced effects for these stimuli. In the case of checkerboards anodal tDCS entirely eliminates the inversion effect, that is, following stimulation there is no significant difference between upright and inverted checkerboards; for faces on the other hand, while anodal stimulation results in a significant reduction in the inversion effect (compared to sham), it does not eliminate it entirely. A significant difference remains between upright and inverted faces under anodal tDCS. There is perhaps an obvious explanation for this stemming from the fact that checkerboards are a more difficult class of stimuli to discriminate prior to any manipulations, Civile, Verbruggen, et al. (2016) found performance for checkerboard to be fairly low (and in some cases to be below chance though not significantly so) and thus performance for the upright stimuli is easier to disrupt than for faces. Civile, Quaglia, et al. (2021) address this by directly comparing the effects of tDCS on the inversion effect for faces and checkerboards using a matching task, discrimination here is much easier than in an old/new recognition task and as such overall performance was high and both the sham inversion effects and performance for the upright stimuli were not significantly different across stimulus types. They used the well-established tDCS montage with anodal stimulation delivered over Fp3 compared to sham and faces and checkerboards were compared in a within-subjects design. The results showed that, in line with previous findings, the inversion effect for checkerboards was eliminated while that for faces was only reduced significantly. Their interpretation of this difference is that while faces and familiar checkerboards share an expertise related component in perceptual learning which can be modulated by the tDCS, face also have an additional component influencing the inversion effect which may be specific to them.

1.8 Introduction to the experiments

The research discussed thus far in this chapter provides a background on how the specificity vs expertise debate has developed, the role of perceptual learning in the inversion effect, and the specific types of information that contribute to it. In this thesis I will report a series of experiments which contribute to the literature on these issues using both behavioural and neuroscientific techniques. In chapter 2 I focused on the disparity between the remaining inversion effects for faces and checkerboards under tDCS and the specific tDCS montage that produces these effects. With this in mind I used a similar paradigm to Civile, Quaglia, et al. (2021) with a matching task to ensure comparable overall performance for faces and checkerboards. I also introduced a new active control condition with the active anodal channel at Fp3 and the return channel at Cz, in addition to the typical Fp3-Fp2 montage. This experiment found results in line with Civile, Qualia et al. (2021) and allows us to consider that there may be a role for both expertise and a face specific component in the face inversion effect. It also provided evidence supporting that it is the specific Fp3-Fp2 montage which produces the reduction in the inversion effect for these stimuli. Chapters 3, 4, and 5, contain a series of behavioural experiments that aimed to assess whether holistic information (as indexed by face contour) may be part of this component that contribute to the inversion effect for faces. Chapter 3 utilised a set of scrambled faces which have been previously demonstrated to produce a robust inversion effect (Civile et al., 2014), a blurring manipulation was applied to the contour of these faces and comparisons made between the inversion effects of the blurred contour and normal contour faces. The findings from this experiment showed that when combined with scrambled faces this contour manipulation was sufficient to reduce the inversion effect compared to control. Chapter 4 applied this to normal (non-scrambled) faces in order to explore whether disruption to contour information alone can produce this effect, the first experiment used an identical blurring manipulation to that reported in chapter 2 and the second experiment tested a novel manipulation which aimed to preserve a greater level of distinctiveness between exemplars. The results here also showed some reduction in the inversion effect (although to a lesser extent than for the scrambled faces), as well as a reduction in overall recognition performance for both upright and inverted faces. Following from research showing that controlling for single-feature orientation can result in a reduction in the inversion effect (Civile, et al., 2016), chapter 5 applied a contour manipulation to faces with half of the features rotated 180° (new Thatcherised faces). Here again we see a reduction in the inversion effect although again it remains significant. Taken together these experiments allow us to conclude that contour information does play a role in face recognition and the inversion effect, although the exact nature of this role is not entirely clear. Chapter 6 returns to the tDCS paradigm and aims to clarify how the anodal Fp3-Fp2 montages produces the effects that it does. This was achieved by applying anodal stimulation and then reversing the polarity (through cathodal stimulation) to investigate what impact this has on the inversion effect. Following a reduction in the inversion effect under anodal stimulation, cathodal stimulation is then able to restore the inversion effect back up to control level. This allowed us to support a perceptual learning account over alternative explanations such as a tDCS induced alteration to scan paths. Finally, chapter 7 summarises the main findings in this thesis and discusses the implications of these in terms of the background literature on both the expertise vs specificity debate and the types of information that contribute to the face inversion effect. I offer suggestions for future research that will extend the work reported in this thesis and further contribute to the face recognition literature.

Chapter 2: Modulating Perceptual Learning for Faces and Checkerboards via tDCS

2.1 Introduction to the experiments

The body of tDCS work in the face and checkerboard recognition literature has explored the modulation of the inversion effect in a wide variety of behavioural paradigms and with a range of tDCS set-ups and procedures. The montage adapted by Civile, Verbruggen et al. (2016) involving anodal stimulation delivered over Fp3 with the return electrode at Fp2 has been shown across a plethora of studies to reduce the inversion effect for both checkerboards and faces (Civile, McLaren et al., 2019; Civile, Waguri et al., 2020; Civile, Cooke et al., 2020; Civile, McLaren et al., 2020; Civile, Quaglia et al., 2021; Civile, McLaren et al., 2021; Civile & McLaren, 2022; Civile, McLaren et al., 2023). In addition to the reduction of the inversion effect, there are two key findings from these previous studies that I will focus on in this chapter. The first is that it is the specific Fp3-Fp2 tDCS montage that produces this result. Civile, McLaren et al. (2018) used an active control with anodal stimulation delivered at rIFG and the return electrode placed on the opposite supraorbital areal at Fp1 (essentially both electrodes placed on the opposite side of the scalp but the same distance from one another). Active control refers to a version of stimulation which matches the active tDCS montage in terms of intensity and duration but is delivered over a different area, with either the anodal electrode, the return electrode or both positioned differently. The rIFG was selected by Civile, McLaren et al. (2018) because previous studies had demonstrated that tDCS delivered over this area was effective in other tasks (e.g. go/no go tasks) (see Cunillera et al., 2014, 2016; Jacobson, Javitt, Lavidor, 2011; Stramaccia et al., 2015) but the effect on perceptual learning tasks had not been investigated. It was found to be not significantly different from sham in its impact on the inversion effect, providing the first evidence that it is the specific Fp3-Fp2 montage rather than an effect of tDCS generally that produces the reduction to the inversion effect (through impairment to the upright stimuli). We can also see from Civile, McLaren et al. (2021) that keeping the same Fp2 return position and moving only the anodal stimulation site (to PO8, chosen based on previous findings that the N170 ERP component and its modulation in response to normal/distorted faces and prototype-defined stimuli is greatest at this channel) also does not result in the same modulation of perceptual learning as the Fp3-Fp2 montage (and was not significantly different from sham), although in this case perceptual learning was investigated in the context of the composite face effect rather than the face inversion effect. Having demonstrated that neither moving the anodal and return electrodes together nor moving only the anodal electrode modulates perceptual learning in the way that Fp3-Fp2 does, the only montage manipulation that remains untested is moving the return electrode while anodal stimulation is delivered at Fp3. I have therefore adopted this variation in experiment 1, using an Fp3-Cz set-up (Cz being another commonly used position for the return electrode in tDCS experiments e.g. Peters et al., 2013; Antal et al., 2004; Antal and Paulus, 2008) with the aim of completing the series of active control variations needed to provide strong evidence that modulation of perceptual learning (as indexed by the inversion effect) is specific to the Fp3-Fp2 tDCS montage.

The other key finding from previous studies relates to the apparent disparity between the remaining inversion effects under anodal tDCS for faces and checkerboards. In the vast majority of checkerboard experiments the inversion effect is entirely eliminated by anodal Fp3-Fp2 tDCS. For faces however, it is typical for a significant inversion effect to remain after anodal tDCS, despite being significantly reduced compared to sham. One possible explanation of this relates to the difference in task difficulty between recognising faces and checkerboards. Civile, Quaglia et al. (2021) addressed this issue by directly comparing the reduction in the face and checkerboard inversion effect using a matching task (a much easier task than the typical old/new recognition task, which allowed comparable performance for both stimulus types). The results fell in line with the previous literature, demonstrating a robust inversion effect for faces and checkerboards in the sham condition and a significantly reduced inversion effect in the anodal condition driven by impaired performance for the upright stimuli. Comparison of these reduced inversion effects for face and checkerboards revealed that the remaining inversion effect for faces was significantly higher than the eliminated inversion effect for checkerboards. In this chapter I follow the same behavioural paradigm as Civile, Quaglia et al. (2021) to ensure high level of performance for both faces and checkerboards, and I will apply the established anodal Fp3-Fp2 procedure and the new active control procedure involving the anodal Fp3-Cz montage. In addition to these anodal conditions, both the Fp3-Fp2 and Fp3-Cz configuration will be used in a sham condition.

The aim this chapter is to extend the face and checkerboard inversion effect literature utilising a new active control condition to provide a clearer characterisation of the impact of tDCS on the inversion effect. The experiment will look at three different tDCS montages, the well-established anodal Fp3-Fp2 condition with anodal stimulation delivered at 1.5mA for 10 minutes, the new anodal Fp3-Cz active control condition which also delivers 1.5mA stimulation for 10 minutes, and a sham condition in which stimulation is delivered for only 30 seconds across a 10-minute period (with this having no observable cortical effect). The sham condition allows comparison of the inversion effects for both active conditions to a baseline as the effect is essentially the same as if no stimulation was delivered, it also helps to preserve a double-blind procedure since all participants feel some level of stimulation and thus cannot discern which experimental condition they are in. Cz was chosen as the active control reference channel for this experiment to give us the same electrode set-up as Ambrus et al. (2011), in their experiment with dot patterns they saw that anodal stimulation in this configuration resulted in impairment to categorisation accuracy and a reduced prototype effect. It therefore seemed important to re-examine this Fp3-Cz montage in the context of the face inversion effect to test whether this might also result in the same reduction of the inversion effect seen in Fp3-Fp2 montage. I would expect the inversion effect for both faces and checkerboards to be reduced in the anodal Fp3-Fp2 condition compared to sham. However, based on the previous literature I would expect the inversion effect for faces to remain significant (despite being significantly reduced) but the inversion effect for checkerboards to no longer be significant, which would confirm Civile, Quaglia et al (2021). The effect of anodal Fp3-Cz condition compared to sham and anodal Fp3-Fp2 is less obviously hypothesised. If anodal Fp3-Cz reduces the inversion effects similarly to anodal Fp3-Fp2, this would provide evidence that the tDCS effects found thus far are likely to be solely related to stimulation of the Fp3 area. However, if it emerges that the effects of anodal Fp3-Cz are similar to sham, this will indicate that the effects found so far in the literature are specific to the montage used (Fp3-Fp2) which involves an anode at Fp3 and return at Fp2.

As a note, the two experiments reported here are currently part of a manuscript submitted for a special issue titled “Advances in Perceptual Learning” to be published in *Journal of Cognitive Enhancement*.

2.2 Methods experiment 1

2.2.1 Participants

In total 120 participants (female=102, mean age= 20.5 years, age range=18-33 years) took part in both experiments (within-subjects). Participants were students from the University of Exeter, recruited largely from the Psychology undergraduate course and pre-screened for right-handedness. Participants were also screened according to tDCS safety criteria to ensure they were safe to participate in the study. All methods were performed in accordance with the relevant guidelines and regulations approved by the Psychology Research Ethics Committee at the University of Exeter

2.2.2 Materials

The face stimuli used in experiment 1a consisted of 256 face images standardised to greyscale on a black background. These were the same stimuli used in Civile, Quaglia et al., 2021 and were originally taken from the Psychological Image Collection at Stirling open database. The faces were cropped such that the hair and neck were removed, leaving a standardised oval shape (see Figure 1). These stimuli were counterbalanced such that each participant saw all of the faces but in a different orientation depending on their group.

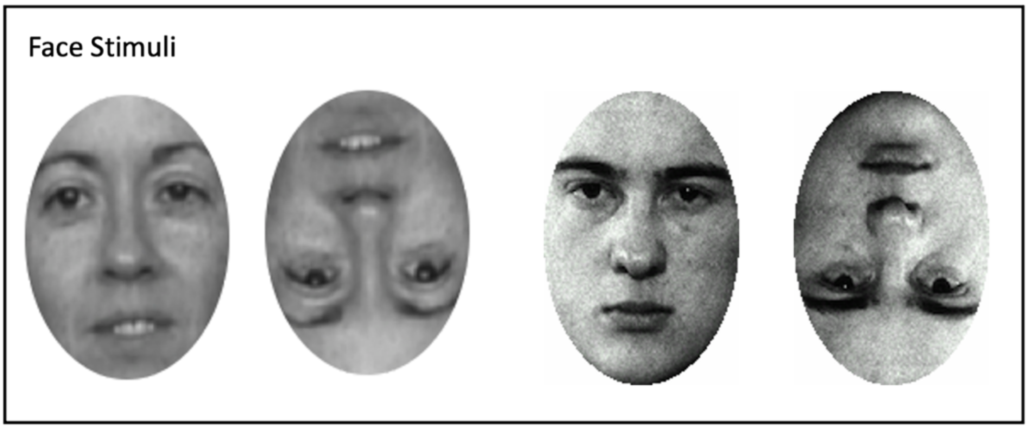

Face Stimuli

Figure 1. Examples of stimuli used in Experiment 1A. Upright and inverted faces standardised to greyscale with the hair and neck cropped to give a uniform oval shape

The checkerboard stimuli in experiment 1b consisted of 16x16 black and white squares. These were the stimuli created by Civile et al. (2014). In each checkerboard 50% of the squares were black and 50% white. Checkerboards were created based on 4 prototype categories, with each exemplar in a given category being generated by altering an average of 48 squares from the original prototype from black to white or vice versa (see Figure 2). The experiments were programmed and run on Superlab 4.0b on an iMac desktop computer with participants positioned around 70cm from the screen

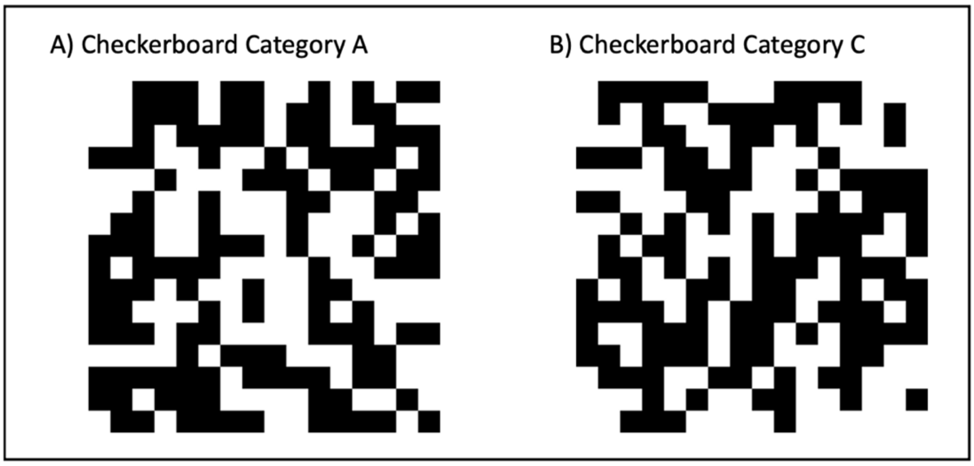

Figure 2. Examples of stimuli used in Experiment 1B. Panel A shows an example derived from prototype category A and panel B shows an example derived from category C.

2.2.3 tDCS apparatus and montage

Participants were split evenly between three tDCS conditions in this experiment. Firstly, the established montage with the active anodal stimulation delivered over Fp3 with a return channel over Fp2, the second was an active anodal condition with stimulation again delivered over Fp3 but the return electrode placed at Cz, and finally a sham condition in which one electrode was always placed over Fp3 and the other was at Fp2 for half of the participants and over Cz for the other half (this was done in order to maintain the double-blind procedure with no difference expected between the two sham set-ups).

The tDCS system used was a neuroConn DC-Stimulator Plus which delivered stimulation through two 35cm^2 electrodes encased in saline soaked sponges. These were placed on the scalp with one at the target area of stimulation and another at the return location and held in place with adjustable head straps. The neuroConn system allowed the use of a double-blind procedure in which a third-party experimenter (not actively running the study) provided numerical codes which were used to determine whether the participant experienced anodal or sham stimulation. Participants in the anodal condition received 10 minutes of stimulation delivered at 1.5mA, while those in the sham condition received 30 seconds of 1.5mA stimulation followed by 0.1mA stimulation delivered for a total of 15 milliseconds spread over the 10-minute period. Each condition began with a 5 second fade-in building intensity up to 1.5mA and ended with a 5 second fade-out reducing the intensity back down to 0. In both conditions, stimulation began as soon as participants started the computer task (see Figure 3).

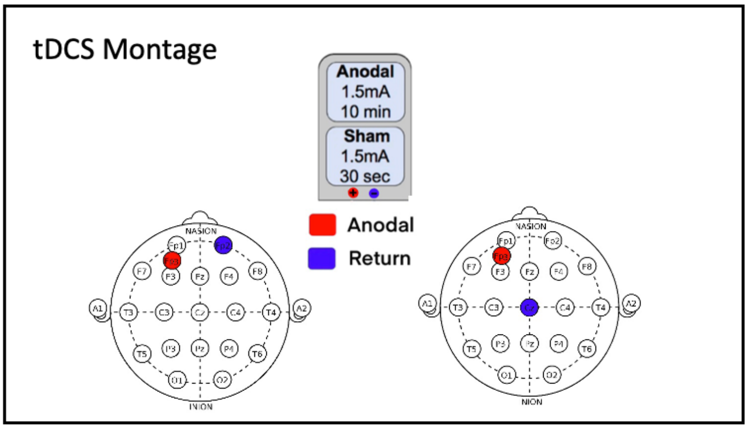

Figure 3. A representation of the tDCS montages used in Experiments 1A and 1B. In both the anode was placed over Fp3. In one set-up the return electrode/cathode was placed over Fp2 and in the other the return electrode/cathode was placed over Cz. Participants were randomly assigned to one of the set-ups and and to active or sham stimulation. The participants from Fp3-Fp2 sham and Fp3-Cz sham were pooled to complete the sham group and there was an equal number of participants in each of the conditions Fp3-Fp2, Fp3-Cz, sham.

2.2.4 Procedure

Participants took part in both the face experiment (1a) and the checkerboard experiment (1b) in a within-subjects design, with the order of these counterbalanced across participants.

Experiment 1a involved a same/different matching task lasting 128 trials. Each trial consisted of a 1s fixation cue presented in the centre of the screen, a 1s presentation of the target face stimulus, a 1.5s mosaic mask, and a test face stimulus shown for a maximum of 2s. Participants were asked to respond using the “X” and “.” keys to indicate whether they thought the test stimulus was “same” or “different” to the target stimulus (the meaning assigned to these keys was counterbalanced across participants). If participants did not respond within 2s they were timed out and the next trial began automatically. Half of the trials contained upright faces and half contained inverted faces (with the target and test stimuli always being the same orientation) and these were presented intermixed and in random order (see Figure 4).

Experiment 1b began with a pre-exposure phase in which participants categorised checkerboards into one of two prototype groups over 128 trials. Each trial began with a 1s fixation cue in the centre of the screen which was followed by a checkerboard presented for up to 4s. Half of the checkerboards were drawn from category A and half from category C, these were presented one at a time, in random order, and participants were asked to respond with which group they thought each checkerboard belong to using the “X” and “.” keys (response keys were counterbalanced across participants). This was done through trial-and-error and after each response participants received feedback indicating whether their response was correct or incorrect, if no response was given after 4s they were timed out and the next trial began (see Figure 4). Following this, participants engaged in a same/different matching task similar to the one described in experiment 1a. For this task half of the stimuli were from category A and half from category C, but none of the exemplars used were those that had been seen in categorisation task previously. The experiment consisted of 128 trials, each consisting of a 1s fixation cue in the centre of the screen, a 1s target checkerboard stimulus, a 1.5s mosaic mask, and a test stimulus shown for a maximum of 2s. Participants were asked to respond using the “X” and “.” keys to indicate whether they thought the test stimulus was “same” or “different” to the target stimulus (the meaning assigned to these keys was again counterbalanced across participants but kept consistent with those used in the face matching task). If participants did not respond within 2s they were timed out and the next trial began automatically. Half of the trials contained upright checkerboards and half contained inverted checkerboards (with the target and test stimuli always being the same orientation) and these were presented intermixed and in random order. Participants were split evenly between three tDCS conditions (anodal Fp3-Fp2, anodal Fp3-Cz, Sham) in this experiment. Stimulation began at the start of the behavioural task, regardless of whether participants engaged with the checkerboard or face experiment first.

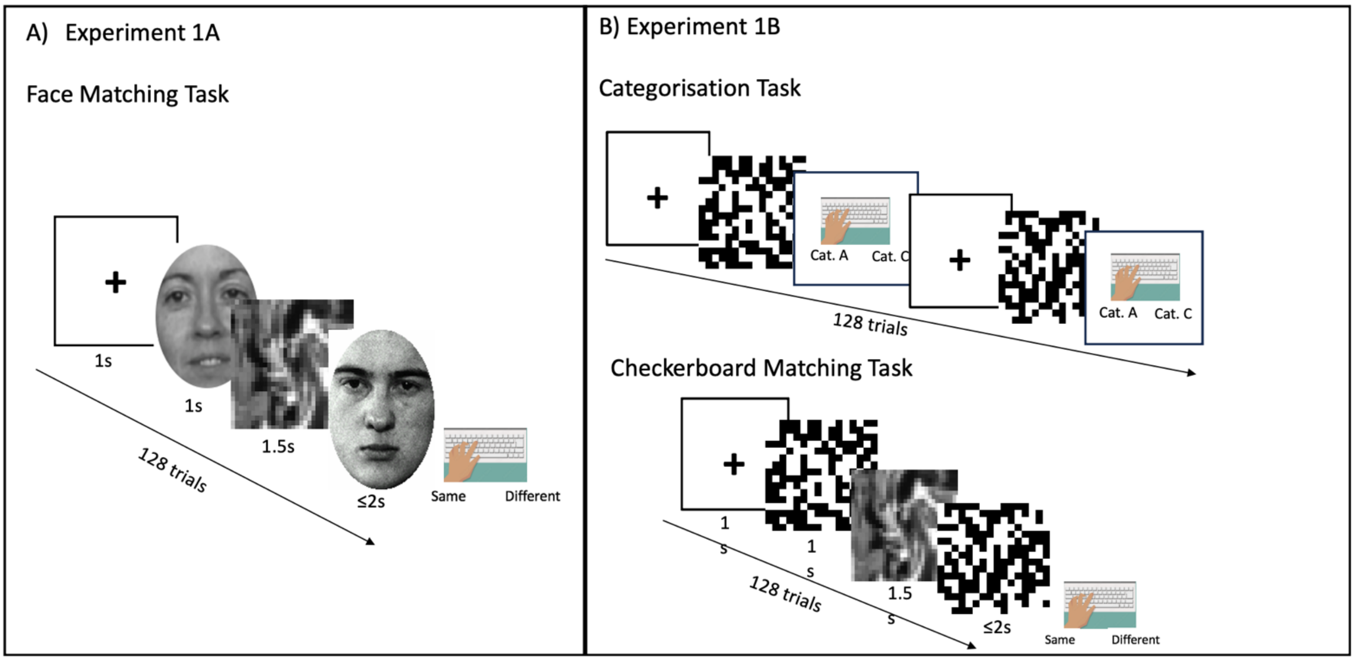

Figure 4. Panel a illustrates the behavioural task adopted in Experiment 1a. The original face images were selected from the Psychological Image Collection at Stirling open database, (http://pics.stir.ac.uk). Panel b illustrates the behavioural task adopted in Experiment 1b. The checkerboard exemplars were selected from Civile et al (2014) and Civile, Quaglia et al (2021).

2.3 Results

2.3.1 Experiment 1a:

A 2 x 3 mixed-model ANOVA with the within-subject factor *face orientation* (upright, inverted) and the between-subjects factor *tDCS condition* (Fp3-Fp2, Fp3-Cz, sham) revealed a significant main effect of *face orientation* *F*(1,117)=86.65, p<.001, η^2p=.42, evidencing the expected face inversion effect, and no significant main effect of *tDCS condition* *F*(1,117)=.46, p=.62, η^2p<.01, indicating that tDCS stimulation does not result in a blanket reduction to performance. Importantly, a significant two-way interaction *face orientation x tDCS condition* *F*(1,117)=5.31, p=.006, η^2p=.08, was also found.

Following this, paired t-tests were conducted on the inversion effect (upright compared to inverted faces) for each tDCS group. In line with previous research the sham group displayed the standard inversion effect with performance for upright faces higher than for inverted faces (M(difference)=.51, SD=.50), *t*(39)=6.52, p<.001, η^2p=.52. In the Fp3-Cz condition a similarly sized inversion effect was found as in the sham condition (M(difference)=.50, SD=.46) *t*(39)=6.96, p<.001, η^2p=.55. In the Fp3-Fp2 condition a reduced inversion effect was found (M(difference)=.20, SD=.48), *t*(39)=2.69, p=.01, η^2p=.15.

Independent t-tests were conducted to compare the size of the inversion effects in each tDCS group. There was no significant difference found in the inversion effect between the sham condition and the Fp3-Cz condition, *t*(78)=.10, p=.91, η^2p<.01, while the inversion effect in the Fp3-Fp2 condition was found to be significantly reduced compared to both the sham condition, *t*(78)=2.80, p=.006, η^2p=.09, and the Fp3-Cz condition, *t*(78)2.82, p=.005, η^2p=.09.

Additionally, performance for upright faces alone was compared across the tDCS groups based on the previous literature demonstrating that the tDCS procedure impacts upright but not inverted faces. No significant difference was found in performance for upright faces in the sham condition compared to the Fp3-Cz condition, *t*(78)2.38, p=.70, η^2p<.01. In the Fp3-Fp2 condition performance for upright faces (M=2.99, SE=.11) was significantly reduced compared to the sham condition (M=3.25, SE=.05), *t*(78)2.03, p=.04, η^2p=.05, while compared to the Fp3-Cz condition (M=3.22, SE=.05) there was no significant difference, *t*(78)=2.03, p=.08, η^2p=.04 (see Figure 5).

2.3.2 Experiment 1b:

A 2x3 mixed-model ANOVA with the within-subject factor *checkerboard orientation* (upright, inverted) and the between-subjects factor *tDCS condition* (Fp3-Fp2, Fp3-Cz, sham) revealed a significant main effect of *checkerboard orientation, F*(1, 117)=12.91, p<.001, η^2p=.09, evidencing the expected checkerboard inversion effect. No significant main effect of *tDCS condition* was found *F*(1, 117)=.21, p=.80, η^2p<.01, again indicating that tDCS stimulation does not result in a blanket reduction to performance. The two-way interaction *checkerboard orientation x tDCS condition* was found to be significant *F*(1, 117)=3.03, p=.05, η^2p=.05.

Following this, paired t-tests were conducted on the inversion effect (upright compared to inverted checkerboards) for each tDCS group. In the sham condition a significant inversion effect was found (M(difference)=.35, SD=.70) *t*(39)=3.15, p=.003, η^2p=.20, and a similarly sized inversion effect was found in the Fp3-Cz condition (M(difference) =.30, SD =.71), *t*(39)=2.65, p=.01, η^2p=.15. The inversion effect in the Fp3-Fp2 condition was eliminated with no difference between upright and inverted stimuli (M(difference) =0, SD=.58), *t*(39)=.09, p=.92, η^2p<.01.

Independent t-tests were conducted to compare the size of the inversion effects in each tDCS group. There was no significant difference found in the inversion effect between the sham condition and the Fp3-Cz condition *t*(78)=.30, p=.76, η^2p<.01, while the inversion effect in the Fp3-Fp2 condition was found to be significantly reduced compared to both the sham condition *t*(78)=2.84, p=.02, η^2p=.06 and the Fp3-Cz condition *t*(78)=2.01, p=.04, η^2p=.09.

Analysis for upright checkerboards alone showed no significant difference in performance between the Fp3-Fp2 condition (M=2.82, SE=.16) and the sham condition (M=3.08, SE=.13) *t*(78)=1.24, p=.21, η^2p=.02, or the Fp3-Cz condition (M=3.08, SE=.12) *t*(78)=1.26, p=.21, η^2p=.02. There was also no difference in upright performance between the sham and Fp3-Cz condition *t*(78)=0, p=.99, η^2p<.01 (see Figure 5).

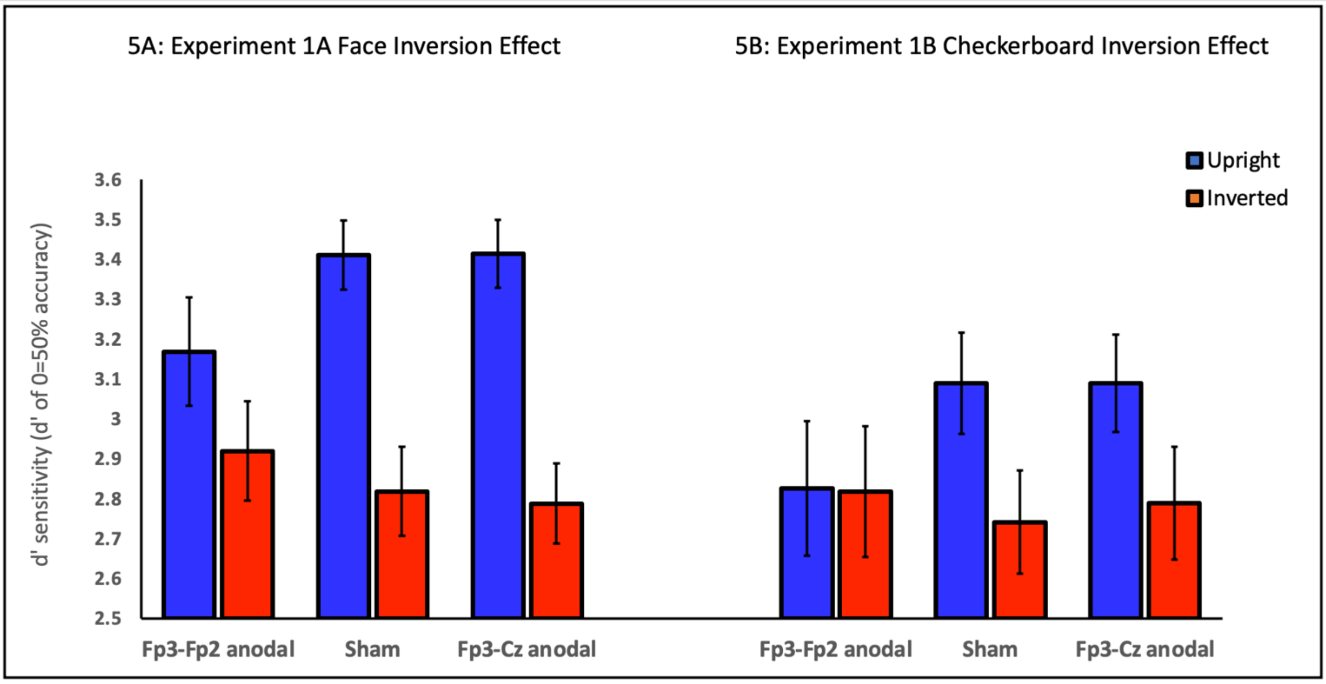

Figure 5. 5a reports the results from Experiment 1a. 5b reports the results from Experiment 1b. The x-axis shows the tDCS conditions, the y-axis shows d’. Error bars represent s.e.m.
In both experiments, performance in all tDCS groups was significantly above chance (for all conditions we found p < .001 for this analysis).

2.3.3 Analyses across experiments

Previous literature has demonstrated that in the Fp3-Fp2 condition the face inversion effect was still significant and larger than the entirely eliminated checkerboard inversion effect (Civile, Quaglia et al., 2021). Analysis comparing the reduced inversion effect for faces in experiment 1a and the eliminated inversion effect for checkerboards in experiment 1b found no significant difference $t(39)=1.67$, p=.10, $\eta^2 p=.07$. Overall recognition performance was also compared for all stimuli averaged together across experiment 1a (M=2.94, SE=.04) and experiment 1b (M=2.89, SE=.08) and no significant difference was found $t(39)=.78$, p=.43, $\eta^2 p<.01$.

2.4 Bayes factor analyses

2.4.1 Experiment 1a

According to the procedure devised by Dienes (2011) we conducted a Bayes analysis on the difference between sham and anodal Fp3-Fp2. We used as the *priors* the differences found in Civile, Quaglia et al. (2021) setting the standard deviation of p (population value | theory) to the mean for the difference between the face inversion effect in sham group vs that in the anodal Fp3-Fp2 group (0.39). We used the standard error (0.07) and mean difference (0.31) between the face inversion effect in the sham group vs. that in the anodal Fp3-Fp2 group. This gave a Bayes factor of 4727, which is very strong evidence (greater than 10, for the conventional cut-offs see Jeffrey, 1961; Dienes 2011) that these results are in line with previous work.

We then conducted a Bayes analysis using as the *priors* the differences found in Civile, Quaglia et al. (2021) between the face inversion effect in sham group vs that in the anodal Fp3-Fp2 group (0.39), but this time we used the standard error (0.07) and mean difference (0.30) between the face inversion effect in the anodal Fp3-Cz group vs. that in the anodal Fp3-Cz group in our Experiment 1a. This gave a Bayes factor of 2586, which is also strong evidence in support that these results are in line with those from the sham vs. anodal Fp3-Fp2 groups.

We conducted a final Bayes analysis using as the *priors* the differences found in Civile, Quaglia et al. (2021) between the face inversion effect in sham group vs that in the anodal Fp3-Fp2 group (0.39), but this time we used the standard error (0.07) and mean difference (0.01) between the face inversion effect in the sham group vs that in the anodal Fp3-Cz group in our Experiment 1a. This gave a Bayes factor of .19 (<1/3), which is moderate evidence for the null supporting the claim that anodal at Fp3-Cz does not impact the inversion effect compared to sham.

Similarly, we also conducted a Bayes factor analysis using as priors the mean difference between sham upright faces and anodal upright faces found in Civile, Quaglia et al (2021) (0.37). We then used the standard error (0.08) and mean difference (0.25) between sham upright vs anodal Fp3-Fp2 upright faces in

Experiment 1a. This gave a Bayes factor of 44, which is also very strong evidence for the position that performance for upright faces is reduced by the anodal Fp3-Fp2 tDCS procedure, consistent with previous results.

Using as priors the mean difference between sham upright faces and anodal upright faces found in Civile, Quaglia et al (2021) (0.37), we then used the standard error (0.09) and mean difference (0.25) between anodal Fp3-Cz vs Fp3-Fp2 upright faces in Experiment 1a. This gave a Bayes factor of 18, which is also strong evidence (greater than 3) for the position that performance for upright faces is reduced by the anodal Fp3-Fp2 montage relative to the Fp3-Cz tDCS procedure.

Finally, using the same priors (0.37), we then used the standard error (0.04) and mean difference (0.02) between sham vs. anodal Fp3-Cz, upright faces in Experiment 1a. This gave a Bayes factor of 0.17, which is moderate evidence for the position that performance for upright faces is not reduced by the anodal Fp3-Cz tDCS procedure.

2.4.2 Experiment 1b

For Experiment 1b we conducted the same Bayes analyses as that for Experiment 1a but this time using as priors the means from Civile, Quaglia et al (2021)'s checkerboard inversion effect. We first set the standard deviation of p (population value | theory) to the mean for the difference between the checkerboard inversion effect in the sham group vs that in the anodal Fp3-Fp2 group (0.57) from Civile, Quaglia et al (2021). We used the standard error (0.10) and mean difference (0.34) between the inversion effect in the sham group vs that in the anodal Fp3-Fp2 anodal group in Experiment 1b. This gave a Bayes factor of 94, which is very strong evidence that these results are in line with our previous work indicating that the Fp3-Fp2 anodal tDCS reduces the checkerboard inversion effect.

Using the same priors (0.57) we then used the standard error (0.09) and mean difference (0.29) between the inversion effect in the anodal Fp3-Cz group vs. that in the anodal Fp3-Fp2 anodal group in Experiment 1b. This gave a Bayes factor of 49, which is very strong evidence that these results are in line with those from the sham vs. anodal Fp3-Fp2 groups.

We conducted a final analysis on the inversion effect differences using the same priors (0.57), but this time we used the standard error (0.11) and mean difference (0.04) between the checkerboard inversion effect in the sham group vs that in the anodal Fp3-Cz group in our Experiment 1b. This gave a Bayes factor of .25, which is strong evidence for the null supporting the claim that anodal at Fp3-Cz does not impact the inversion effect compared to sham.

Similarly, we also conducted a Bayes factor analysis using as priors the mean difference between sham upright checkerboards and anodal upright checkerboards found in Civile, Quaglia et al (2021) (0.43). We then used the standard error (0.13) and mean difference (0.26) between sham upright vs anodal Fp3-Fp2 upright checkerboards in Experiment 1b. This gave a Bayes factor of 3.5, which is moderate evidence for the position that performance for upright checkerboards is reduced by the anodal Fp3-Fp2 tDCS procedure, consistent with previous results.

With the same priors (0.43), we then used the standard error (0.11) and mean difference (0.26) between anodal Fp3-Cz vs Fp3-Fp2 checkerboards faces in Experiment 1b. This gave a Bayes factor of 6.7, which is also moderate evidence for the position that performance for upright checkerboards is reduced by the anodal Fp3-Fp2 montage relative to the Fp3-Cz tDCS procedure.

Finally, using the same priors (0.43), we then used the standard error (0.13) and mean difference (0.0) between sham vs. anodal Fp3-Cz, upright checkerboards in

Experiment 1b. This gave a Bayes factor of 0.28, which is strong evidence for the null, supporting the position that performance for checkerboards is not reduced by the anodal Fp3-Cz tDCS procedure.

2.4.3 Additional Analysis

We conducted a final Bayes analysis using as priors the difference between the face inversion effect and the checkerboard inversion effect from the anodal Fp3-Fp2 groups in Civile, Quaglia et al (2021) (.43), using the standard error (0.09) and mean difference (0.20) for the same difference in our experiments. This gave a Bayes factor of 4.3 which is moderate evidence that our results are in line with previous ones supporting how the remaining face inversion effect in the Fp3-Fp2 group is larger than the fully reduced checkerboard inversion effect in the same group.

2.5 Discussion

In these experiments I have expanded on the body of research exploring how the tDCS montage adapted by Civile, Verbruggen et al. (2016) can modulate the inversion effect as an index of perceptual learning. I have demonstrated that anodal stimulation delivered over Fp3 with a return electrode at Fp2 significantly reduces the inversion effect for both faces (experiment 1a) and checkerboards (experiment 1b) which is consistent with the previous literature (Civile, Verbruggen et al., 2016; Civile, McLaren et al., 2018; Civile, Obhi et al., 2019; Civile, McLaren et al., 2020; Civile, McLaren et al., 2021; Civile, Quaglia et al., 2021; Civile & McLaren 2022; Civile, McLaren et al., 2023). In addition, both experiments 1a and 1b have provided new evidence from the active control condition. From previous work we can see that active stimulation delivered over different anodal and return sites. (rIFG-Fp1) produces no reduction in the inversion effect compared to sham (Civile, McLaren et al., 2018), the same is observed when only the anodal electrode is moved (PO8-

Fp2) (Civile, McLaren et al., 2021). In these experiments we have also demonstrated that in an active control where only the return electrode is moved (Fp3-Cz) there is no reduction in the inversion effect compared to sham for either faces or checkerboards. This has helped extend our understanding of how tDCS can modulate perceptual learning by providing strong evidence that it is only the specific Fp3-Fp2 tDCS montage that results in a significant reduction in the inversion effect, and even when Fp3 is targeted with anodal stimulation no effect is found with a different return site. The results also confirm that there is no effect on overall performance under anodal tDCS for either faces or checkerboards but that the Fp3-Fp2 montage specifically disrupts perceptual learning as indexed by the inversion effect.

Furthermore, in experiment 1a I have again demonstrated that the reduction in the face inversion effect in the Fp3-Fp2 condition compared to sham or Fp3-Cz is driven by impaired performance for upright faces. This finding was supported by the Bayes factor analysis conducted between this experiment and Civile, Quaglia et al., 2021 in which it was confirmed that it is performance for upright faces that is impacted by the tDCS. For the checkerboards however we found only a numerical reduction in performance for upright checkerboards in the Fp3-Fp2 condition compared to sham and Fp3-Cz with no statistically significant difference. The Bayes factor analysis however did provide some support for the notion that, as in Civile, Quaglia et al., 2021, the upright checkerboards were disrupted by the tDCS montage. Additionally, while our initial analysis found no significant difference between the reduced inversion effect for faces and the eliminated inversion effect for checkerboards, the Bayes factor analysis found robust evidence that the findings were in line with the previous research from Civile, Quaglia et al., 2021 which showed that the face inversion effect remained higher than the checkerboard inversion effect.

The conceivable disparity between the face Inversion effect and the checkerboard inversion effect under the Fp3-Fp2 tDCS montage may have a number of potential explanations; it could be that we simply have vastly more expertise in recognising faces given our life-time of experience and as such this level of perceptual learning is much harder to eradicate than the relatively small amount of expertise participants gained for checkerboards in the pre-exposure phase. Another explanation may be that both faces and checkerboards share perceptual expertise as one factor that contributes to the inversion effect and which is disrupted by tDCS, but that for faces there is an additional (possibly face-specific) mechanism which is unaffected by the tDCS which accounts for the remaining inversion effect. Given the inherent impracticality of matching level of expertise for these two types of stimuli, in the following chapters I will follow the second line of reasoning in a series of behavioural experiments to explore which types of information might be contributing to the face inversion effect but remain unaffected by the tDCS.

Chapter 3 Manipulating Face Contour for Scrambled Faces

3.1 Introduction to the experiments

In the next series of chapters I explore the idea that there may be some aspects of the face stimuli which have not been previously manipulated in the inversion effect literature and aim to assess whether these types of information may contribute to both overall recognition performance and the inversion effect. Findings that support this may indicate that there is face-specific information contributing to the inversion effect (in addition to perceptual expertise) which may account for the remaining inversion effect observed in experiment 1 in the Fp3-Fp2 condition.

The roles that single feature orientation, first-order configural information and second-order configural information play in the inversion effect have been investigated extensively in the previous literature although not always with consistent results. Leder and Bruce (1998) manipulated the individual features or configural information in face stimuli which had previously been rated as average in distinctiveness (resulting in three versions of each face: the original, the featural manipulation, and the configural manipulation). The manipulations artificially increased how distinctive the faces were, and allowed direct comparison to show which type of information was most impactful. When faces were presented upright, both the featural and configural manipulations resulted in participants rating the distinctiveness as higher compared to the original faces. However, when the faces were inverted the distinctiveness of those with the configural manipulation was more greatly reduced than those with the featural manipulation. This pattern of results also held true for recognition performance, indicating that configural information may be more susceptible to disruption by inversion than other types of information and therefore be contributory to the face inversion effect. Evidence for the role of second-order configural information in the inversion effect can be demonstrated in comparisons of recognition performance for facial features in their usual configuration and when the positioning of the eyes is manipulated (Tanaka & Sengco, 1997). Recognition performance for features in their usual configuration was higher than that for the disrupted configuration when they were presented upright but showed no difference in performance when inverted. In terms of the inversion effect this demonstrates how second-order configural information may confer a benefit for upright faces but is not utilised in the same way for inverted faces.

The role of featural information has also been investigated and in light of this previous evidence it seems logical that when conducting meta-analytic research into the role of featural information, McKone and Yovel (2009) initially hypothesised that inversion would have minimal impact on sensitivity to featural information. This was not however what their findings revealed, rather their analysis of 22 inversion effect studies indicated that the orientation of the individual features also contributes to the inversion effect and may do so to an equal degree as configural information.

Perhaps the most relevant studies for the current work are those conducted on face stimuli by Civile et al. (2014) and Civile et al. (2016) which aimed to manipulate the types of information available to participants. In Civile et al. (2014) a set of scrambled faces was generated by selecting at random one of main the facial features (the eyes individually, the nose, the mouth and the ears individually) and repositioning them starting from the forehead and moving the others in sequence so that each feature takes the place of the last. In this way they were able to disrupt first-order configural information by changing the spatial relationship the features have with one another, and second-order configural information by changing that relationship in comparison to a prototypical face. The specific scrambling manipulation used here aimed to address some of the issues present in previous literature using scrambled faces. They give as an example Collishaw and Hole (2000) who always moved the eyes up or down the faces as a unit and also moved the ears and nose together, preserving the first-order and second-order configural relationship between these features. Civile et al. instead move each feature individually, with the positions arranged such that the configural information is disrupted entirely. In addition, the scrambled faces were designed to follow a prototypical configuration, just as normal faces do. They created four different prototype configurations with the sequence of features moved different in each, these were then counterbalanced across participants groups to ensure that their findings were not the result of a given configuration being particularly easy or difficult for participants to recognise in the upright orientation. In an old/new recognition task participants were presented these stimuli both upright and inverted along with normal faces and the findings showed that this disruption to the configural information reduced overall recognition performance but did not reduce significantly the inversion effect.

These findings were substantiated by Civile et al. (2016) who as part of their series of studies ran a full replication of the old/new recognition task with the scrambled faces and again found a robust inversion effect equivalent to that found for normal faces. From these we can see that disruption to first and second-order configural information is not sufficient to significantly reduce the face inversion effect. In Civile et al. (2016)'s interpretation of the inversion effect for scrambled faces they suggest that holistic information may be one of the thus far untested factors contributing to the inversion effect. This follows from the theory presented by Hole, George and Dunsmore (1999) which makes a distinction between two different types of relational processing, configural and holistic. They argue that evidence for this distinction can be seen in their discovery that photographic negatives of faces are subject to some configural processing (demonstrated in their experiment on the chimeric face effect) despite research showing that configural disruptions such as Thatcherisation are harder to detect in negatives compared to positives (Lewis & Johnston, 1997). Additionally, the chimeric face effect has been observed to be comparable in young children and adults in young children, while the inversion effect grows more pronounced with age (Carey & Diamond, 1994). Hole et al.'s (1999) theory therefore states that configural processing is a fine-tuned mechanism (which likely develops with experience, and therefore age) related to the specific positioning of the facial features (which is heavily disrupted by the scrambling manipulation), while holistic processing is a broader mechanism which is elicited by stimuli that follow the rough plan of a face and allows you to identify a stimulus as a face. Following this line of reasoning I decided to design a series of studies focusing on holistic information and chose the contour/outline of the face to use as a manipulable component of this. The advantage of selecting the face outline to manipulate is that it does not affect the configural information amongst the main internal facial features e.g., the distance between the eyes, ears, nose, and mouth remains unaffected, as does their configuration. Importantly, the face outline presented alone (i.e., without any of these internal features) would still clearly represent the rough plan of a face referred to in the theory laid forth by George and Dunsmore (1999). The following study extends the scrambled faces paradigm to include stimuli for which the face contour has also been altered with the aim of investigating whether this type of information may contribute to the face inversion effect.

3.2 Methods experiment 2

3.2.1 Participants

In total, 144 participants (female=89, mean age=21.7, age range=16-57) took part in this experiment. 72 of these were students from the University of Exeter, and 72 were recruited through the third-party recruitment service Prolific. Both groups were compensated for their participation. Analyses with *Recruitment* as a factor (University or Prolific) showed no main effect ($F(1, 142)=.262$, $p=.60$, $\eta^2_p<.01$) and it did not interact significantly with any other factors in the study (max. $F(1, 142)=2.57$, $p=0.11$, $\eta^2 p= 0.018$). The sample size was determined from previous studies utilising the same face stimuli, counterbalancing of the participant conditions and stimuli, and behavioural paradigm (Civile et al., 2014; Civile et al., 2016).

3.2.2 Materials

This study used two sets of stimuli, both based on the sets of scrambled faces adopted in Civile et al (2014) and Civile et al (2016). These consisted of 128 male faces, standardised to a greyscale colour on a black background, each with a neutral expression and the hair and neck cropped out (leaving the contour and ears intact). The scrambling manipulation was based on four different prototype groups, each of which had a different configuration of features. To create the scrambled faces, six facial features (the mouth, nose, two ears, and two eyes (including eyebrows)) were rearranged in specific order depending on the prototype group the exemplar is drawn from. In each case a feature is chosen at random and moved to the forehead (as this is the largest area of the face that can accommodate a feature) and a second feature is then selected and moved into the space left by the first feature e.g., in category A the right eye moves to the forehead and the right ear then fills the space where the right eye was originally. This sequence of selecting and moving a feature continued until all six facial features had been moved, and the configuration of faces in a given prototype category is consistent with all others in that group e.g., all faces in category A will have the right eye on the forehead, the right ear in the right eye space etc. These groups of scrambled faces completed one set of stimuli; the second set began with these scrambled faces but were additionally altered by blurring the contour of the face outward such that the outline is no longer a recognisable face shape (see Figure 6). As before, stimuli were counterbalanced across participants groups meaning each face identity was shown to some participant in every condition, and each participant saw each face identity in some condition, but in this case the four prototype categories were also counterbalanced across participant groups. The experiment was programmed and run on Gorilla.sc and participants used their own laptop or desktop devices.

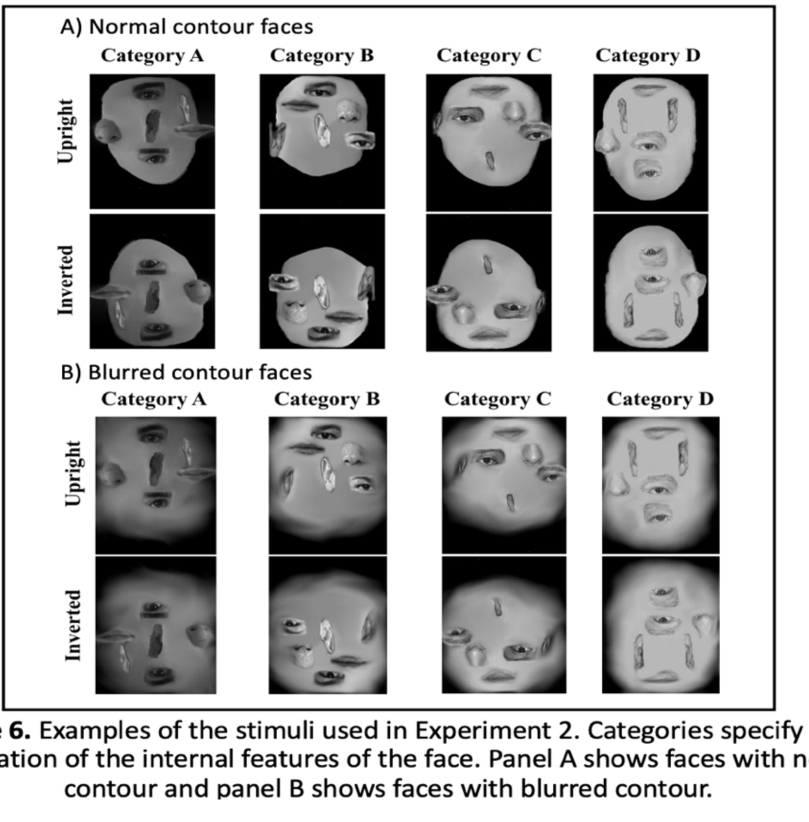

Figure 6. Examples of the stimuli used in Experiment 2. Categories specify the configuration of the internal features of the face. Panel A shows faces with normal contour and panel B shows faces with blurred contour.

3.2.3 Procedure

The behavioural task used was an old/new recognition task consisting of a study phase and a recognition phase. The study phase was run over 64 trials, each one began with a 1s fixation cue in the centre of the screen followed by a scrambled face image presented for 3s. The faces were split evenly between normal-contour and blurred-contour faces with half of each set presented upright and half inverted, these were presented intermixed, in random order. No response was required from participants during the study phase, and they were asked to memorise as many of the faces as possible. The recognition phase consisted of 128 trials with 50% of those involving the stimuli from the study phase and 50% involving novel stimuli presented one at a time in random order. Each trial began with a 1s fixation cue in the centre of the screen, followed by a face stimulus shown for a maximum of 3s. Participants responded using the “X” and “.” keys to indicate whether or not they thought a given stimulus had been shown in the study phase (the meaning assigned to the keys was counterbalanced across participant groups). If no response was given after 3s participants were timed out and the next trial began automatically (see Figure 7).

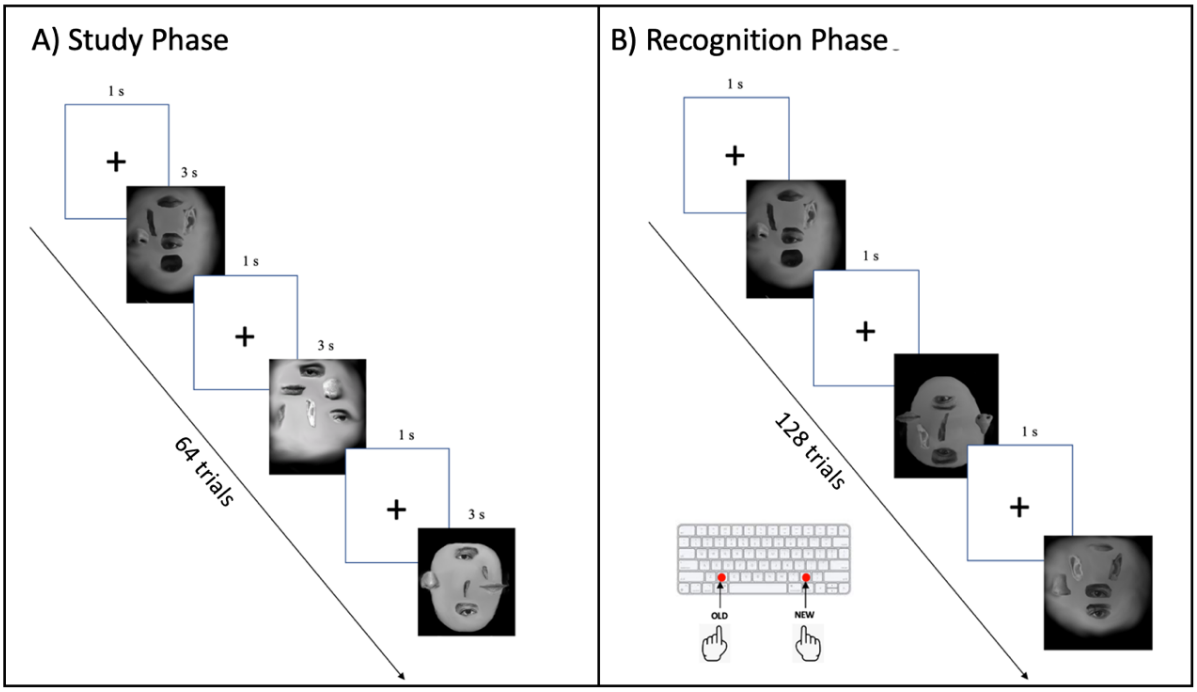

Figure 7. A representation of the old/new recognition task used in Experiment 2. Panel a illustrates the study phase. Panel b illustrates the recognition phase. The recognition phase followed immediately from the study phase, with only a brief break to deliver instructions to participants on how they should respond.

3.3.1 Results

To be consistent with previous related studies (Civile et al., 2014a, 2014b, 2016a), the accuracy scores collected in this experiment were converted to a d-prime (d’) sensitivity measure (Stanislaw & Todorov, 1999) and this is the measure used in my analyses. This was an important step as the old/new recognition paradigm is a yes/no task involving signal and noise trials and d’ sensitivity considers both the hit rate (H) and the false alarm rate (F) and is therefore a more useful measure than the number of correct responses alone. H is calculated based on the proportion of trials to which the correct answer was YES, and the participants responded YES. F is calculated based on the proportion of trials to which the correct answer was NO, and the participants responded YES. d’ sensitivity is a measure of the difference between these, calculated as the difference between the z transforms of the two rates: d’=z(H)-z(F).

We assessed performance against chance (*d'* of 0) which showed that both types of upright faces (normal contour and blurred-contour) were recognized significantly above chance ($p < .001$). Performance against chance for inverted blurred-contour scrambled faces showed a trend towards significance ($p = .052$) whereas just like in Civile et al (2014) and Civile et al (2016) inverted scrambled were not recognized significantly above chance ($p = .45$). When all stimulus conditions are collapsed overall performance is significantly above chance ($p<.001$).

A 2x2 within subjects ANOVA using as factors *Face Type* (scrambled, scrambled blurred-contour) and *Orientation* (upright, inverted) revealed a significant interaction between these factors, $F(1, 143) = 4.41$, $p = .037$, $\eta^2_p = .030$. A significant main effect of *Orientation* was found (upright better), $F(1, 143) = 17.28$, $p < .001$, $\eta^2_p = .108$. No significant main effect for *Face Type* was found, $F(1, 143) = .63$, $p = .425$,

$\eta^2_p < .01$. As in Civile et al (2014) and Civile et al (2016) follow up, paired samples *t*-tests were conducted to compare performance on upright and inverted faces i.e., the face inversion effect, for each face type. We found a large inversion effect for scrambled faces with performance for upright (*M* =.31 *SD* =.48) significantly better than that for inverted scrambled faces (*M* = .04 *SD* = .53), *t*(143) = 4.71, p < .001, η^2_p = .036. Although performance for upright scrambled blurred-contour faces (*M* = .19 *SD* = .57) was numerically higher than that for inverted ones (*M* = .09 *SD* = .51) no significant inversion effect was found, *t*(143) = 1.62, p = .107, η^2_p = .239. The significant interaction can be interpreted as being due to a reduced inversion effect in the scrambled blurred-contour faces.

Importantly, in similar fashion to Civile et al (2014) and Civile et al (2016), we directly compared performance for upright scrambled faces vs that for upright scrambled blurred-contour faces and for inverted scrambled faces vs inverted scrambled blurred-contour faces. These comparisons are particularly appropriate because the same stimulus sets are rotated across participants in a counterbalanced manner; so that for each upright or inverted face seen in a scrambled condition for a given participant will equally often serve as an upright or inverted face in the scrambled blurred-contour condition. Performance for upright scrambled faces was significantly higher than that for scrambled blurred-contour faces, *t*(31) = 2.07, p = .040, η^2_p = .029. No significant difference was found between inverted scrambled and inverted scrambled blurred-contour faces, *t*(31) = .846, p = .399, $\eta^2_p < .01$.

3.3.2 Additional analysis

In light of the potential issues identified around overall performance for the inverted faces being not significantly above chance, I now explore some additional analysis excluding any participant whose overall recognition performance (for all stimuli) was below chance. For experiment 2 this involved the exclusion of 39 participants (of these 18 were recruited through the university and 21 through prolific) leaving a sample size of 105. These exclusions of course brought performance to significantly above chance (with comparisons giving $p<.001$ for normal contour upright, normal contour inverted and blurred contour upright faces, and $p=.003$ for blurred contour inverted faces).

A 2x2 within-subjects ANOVA using the factors *Face Type* (scrambled normal contour, scrambled blurred-contour) and *Orientation* (upright, inverted) revealed a significant main effect of *Orientation*, $F(1, 104)=15.71$, $p<.001$, $\eta^2_p=.131$, demonstrating the robust inversion effect. No significant main effect of *Face Type*, $F(1, 104)=.34$, $p=.346$, $\eta^2_p<.01$, was found showing no difference in overall recognition performance for scrambled faces with normal contour compared to blurred contour. In line with the original analysis, there was a significant interaction (*Face Type x Orientation*) found, $F(1, 104)=4.61$, $p=.034$, $\eta^2_p=.042$. Breaking down this interaction, two paired-samples *t*-tests were conducted, the first comparing upright and inverted performance for the scrambled faces with normal contour. A significant inversion effect was found with upright faces (M=.43 SD=.45) recognised better than inverted (M = .13 SD = .44), $t(104) = 4.95$, $p < .001$, $\eta^2_p= .147$. The second was conducted on the blurred contour faces and although performance for upright faces was numerically higher (M=.28 SD=.57) than for inverted faces (M=.17 SD=.47), there was no significant inversion effect found $t(104)=1.4$, $p=.164$, $\eta^2_p=.01$. Additional *t*-tests comparing the upright and inverted stimuli for each face type found that performance for upright scrambled normal-contour faces was significantly better than that for scrambled blurred-contour faces, $t(104)=2.07$, $p=.042$, $\eta^2_p=.029$, while there was no significant difference found between inverted scrambled normal-contour and inverted scrambled blurred-contour faces, *t*(104)=.757, p=.45, η^2_p<.01 (see Figure 8).

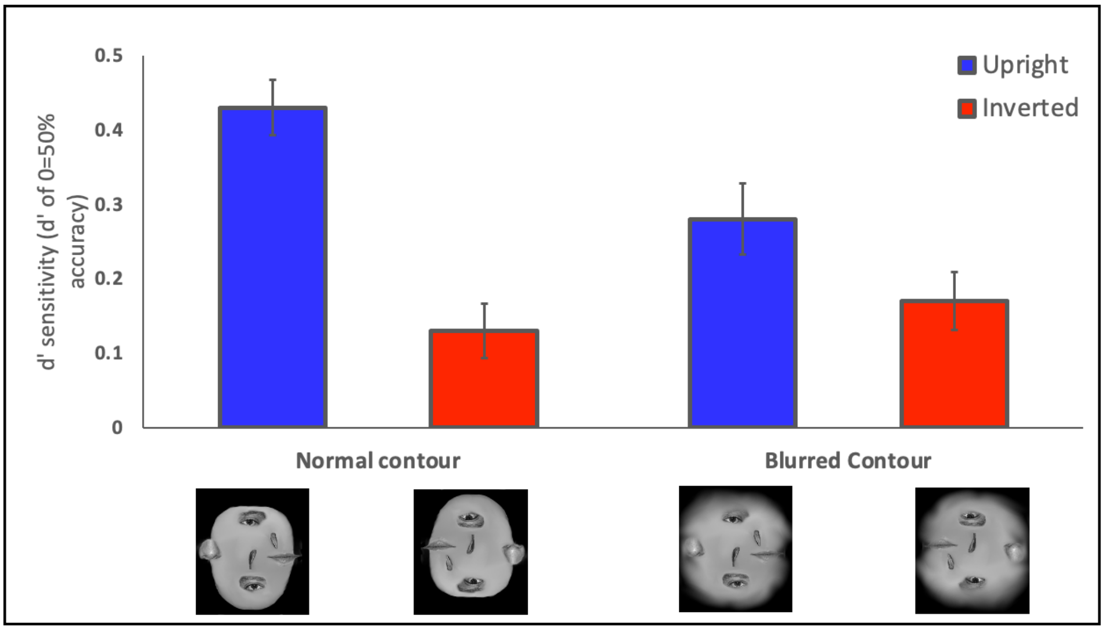

Figure 8. Graph reports the results from Experiment 2 when participants whose performance was below chance were excluded. The x-axis shows the stimulus conditions, scrambled faces with either normal or blurred contour. The y-axis shows d'. Error bars represent s.e.m.

3.4 Discussion

The results of this study firstly confirm the findings from Civile et al. (2014a) and (2016a) that a large inversion effect is still observed for faces when the configural information is disrupted through the scrambling manipulation, when they are derived from prototype defined categories. Given the total disruption to first and second-order that occurs in scrambled faces, these findings indicate that it is likely that these types of information contribute directly to the inversion effect (although they may also factor into overall recognition performance).That the inversion effect in the current experiment persists for the normal-contour scrambled faces supports the previous findings that neither first nor second-order configural information are necessary to obtain an inversion effect, however we do not have further evidence at this point about the effect of scrambling alone on overall recognition performance (as there were no normal face stimuli with which to compare). Chapters 4 and 5 will offer pilot data with this comparison of scrambled face stimuli to normal face stimuli (in addition to a contour manipulation) and expansion on Civile et al. (2014)'s finding that controlling for single-feature orientation may impact the inversion effect.

The novel finding in this research is that disruption to the face contour through the blurring manipulation was sufficient to significantly reduce the inversion effect for these stimuli vs control. In addition, the inversion effect for blurred-contour faces was found to be not significant (although still numerically present). That this finding persists in the additional analysis where participants performing below chance were excluded indicates that this result is not due to rather low overall performance (although significantly above chance) for example resulting in the inverted stimuli hitting a floor effect in the original analysis. This finding is particularly interesting for the account of holistic processing proposed by Hole et al. (1999), this theory makes the distinction between configural processing and holistic processing. The former they suggest is a specialised mechanism that develops with experience and focuses on specific facial details, with the latter being a much broader process which identifies that a stimulus is a face. Hole et al. (1999) posited that holistic processing can be elicited by stimuli that conform to the basic plan of a face, and as such it was thought that contour information may play an important role. Civile et al. (2016)'s work highlighted first-order configural information as another aspect that might contribute to holistic processing, firstly due to their finding that an inversion effect can be obtained with first-order configural information when single-feature orientation is controlled for, and secondly because first-order configural information is also largely common across faces and should therefore help to identify that a given stimulus is a face. The reduction of the inversion effect for the blurred-contour scrambled faces supports the idea that contour information may elicit a type of holistic processing that differs in purpose from the configural processing, and that this type of processing may be causal to the inversion effect. While we do not have direct evidence of holistic processing in this experiment, we can draw the conclusion from this and previous research that the inversion effect can be reduced by altering single-feature or contour information, particularly when the typically familiar configural information is not available to be utilised due to the scrambling manipulation. Taken together the findings from this experiment support the idea the face contour, is a contributary factor in the face inversion effect.

In Chapter 4 I extended the contour manipulation to normal faces. The aim was to investigate how important the face contour is to the inversion effect when all the configural information amongst the main facial faces (eyes, mouth, nose, and ears) are unaltered.

Chapter 4 Manipulating the Face Contour of Normal faces

4.1 Introduction to the experiments

Having shown in experiment 2 that the inversion effect can be greatly reduced when the face contour is manipulated in combination with the scrambling manipulation, I now explore the extent to which face contour alone impacts the inversion effect on normal faces. In Civile et al. (2014a) the scrambling manipulation used was compared directly to normal faces in the old/new recognition task and while the inversion effect was not significantly reduced for scrambled faces compared to normal there was a reduction in overall recognition performance. It may be that participants' sensitivity to the contour manipulation is due in part to recognition for the faces being more difficult than usual at baseline, perhaps configural information amongst the main features is a key component in how faces are usually recognised but in the absence of this being available, participants were forced to rely more on the face contour which may be more sensitive to disruption by inversion. If the findings of experiment 2 persist with normal faces it would provide some evidence that the face contour could be itself a contributory factor in the inversion effect.

This chapter details two experiments following the same behavioural paradigm described in the previous chapter, experiment 3a also employs the same blurring manipulation used in experiment 2 on an otherwise unmanipulated normal face (i.e., with no additional disruption to featural, first-order and second-order information). Following the completion of this study concerns were identified regarding the blurring technique relating to the integrity of the image and the lack of objecthood remaining for the blurred stimuli. In experiment 3b I aimed to address these potential issues and developed an alternative manipulation that involved replacing the contour of the face with a new outline which was based on the original spatial relationship between the internal features of the face and the outline, in this way I was able to preserve the amount of information available in the faces that made them distinguishable from one another while still disrupting the holistic information as indexed by face contour.

4.2 Methods experiment 3a

4.2.1 Participants

In total, 144 participants (female=98, mean age=23.1, age range=18-55) took part in experiment 3a. 72 of these were students from the University of Exeter, and 72 we recruited through the third-party recruitment service Prolific. Both groups were compensated for their participation. Analyses with *Recruitment* as a factor (University or Prolific) showed no main effect ($F(1, 142)=1.59$, p=0.22, $\eta^2 p<.01$) and it did not interact significantly with any other factors in the study (max. $F(1, 142)=2.82$, p=0.10, $\eta^2 p= 0.02$). The sample size was determined from previous studies utilising the same face stimuli, counterbalancing of the participant conditions and stimuli, and behavioural paradigm (Civile et al., 2014; Civile et al., 2016).

4.2.2 Materials

The face stimuli used in this experiment consisted of 128 face images that were standardized to greyscale on a black background. The original face images were selected from the Psychological Image Collection at Stirling open database. One set of faces used in the experiment were male faces with a neutral expression and the hair and neck cropped out leaving only the outline of the face and the internal features. These stimuli were also those used in Civile et al. (2014a) and Civile et al. (2016a). The second set began as these original faces but were additionally altered by blurring the contour of the face outward such that the outline is no longer a recognisable face shape (see Figure 9). Stimuli were counterbalanced across participants groups meaning each face identity was shown to every participant, but the orientation and contour were dependent on the participant group. The experiment was programmed and run on Gorilla.sc and participants used their own laptop or desktop devices.

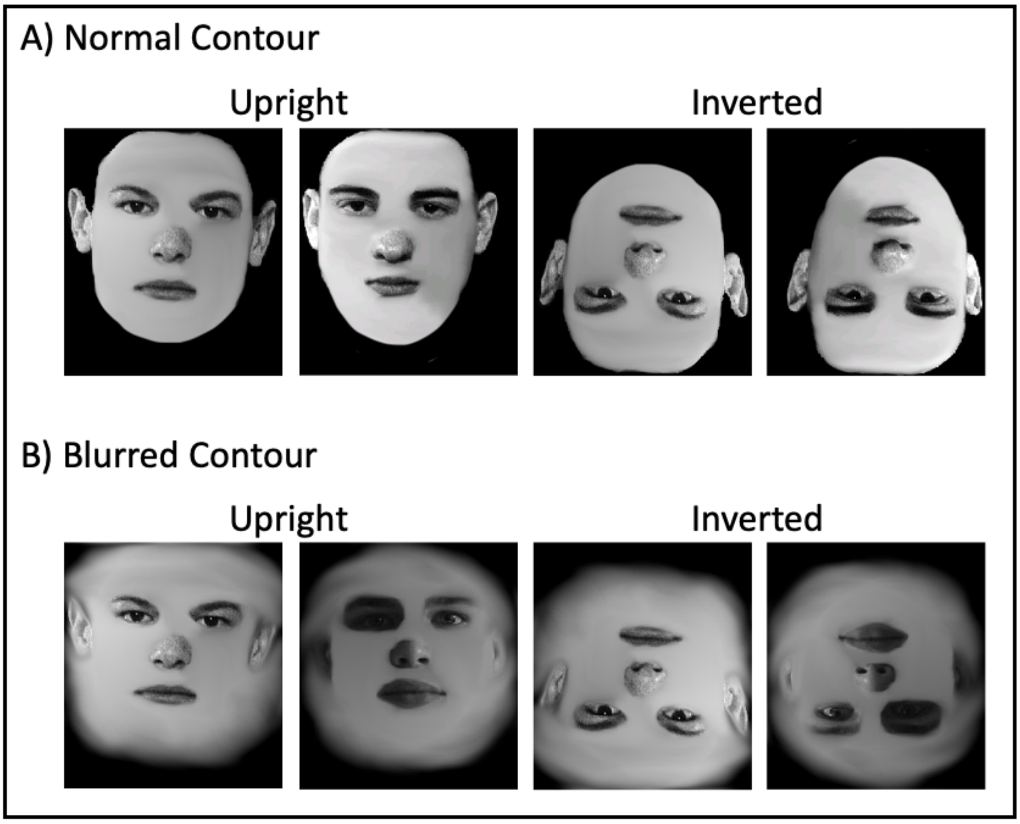

Figure 9. Examples of the stimuli used in Experiment 3a. Panel A shows faces with normal contour and panel B shows faces with blurred contour.

4.2.3 Procedure

In line with experiment 2 the behavioural task used was an old/new recognition task consisting of a study phase and a recognition phase. The study phase was run over 64 trials, each one began a 1s fixation cue in the centre of the screen followed by a face image presented for 3s. The faces were split evenly between normal-contour and blurred-contour faces with half of each set presented upright and half inverted (16 upright normal-contour, 16 inverted normal-contour, 16 upright blurred-contour, 16 inverted blurred-contour), these were presented intermixed and in random order. No response was required from participants during the study phase, and they were asked to memorise as many of the faces as possible. The recognition phase consisted of 128 trials, with 50% of those involving the stimuli from the study phase and 50% involving novel stimuli (also evenly split between normal and blurred contour and upright and inverted orientations presented one at a time in random order). Each trial began with a 1s fixation cue in the centre of the screen, followed by a face stimulus shown for a maximum of 3s. Participants responded using the “X” and “.” keys to indicate whether or not they thought a given stimulus had been shown in the study phase (the meaning assigned to the keys was counterbalanced across participant groups). If no response was given after 3s participants were timed out and the next trial began automatically.

4.3 Results

As before, accuracy scores were converted to a d’ sensitivity measure prior to analysis.

A 2x2 within-subjects ANOVA using the factors *Face Type* (normal contour, blurred-contour) and *Orientation* (upright, inverted) revealed a significant main effect of *Orientation*, $F(1,143)=48.65$, $p<.001$, $\eta^2 p=0.25$, demonstrating the well-established face inversion effect. A significant main effect of *Face Type*, $F(1, 143)=4.57$, $p=0.034$, $\eta^2 p=0.03$, was also found showing that overall recognition performance for normal-contour faces (M=0.49, SD=0.46) was significantly higher than that for blurred-contour faces (M=0.39, SD=0.45). They key finding in this experiment was that no significant interaction (*Face Type x Orientation*) was found, $F(1, 143)=0.90$, $p=0.34$, $\eta^2 p<0.01$, indicating that both sets of faces showed a robust inversion effect, although numerically the inversion effect for normal faces with blurred-contour (M = 0.29, SD = 0.74) was smaller than the one found for normal faces with normal contour (M = 0.37, SD = 0.74).

Importantly, comparison of stimuli against chance showed that performance for upright normal-contour faces ($p<.001$), upright blurred-contour faces ($p<.001$), inverted normal-contour faces($p<.001$), and inverted blurred-contour faces ($p<.001$), were all significantly above chance. When all stimulus conditions were collapsed, overall performance was also found to be significantly above chance ($p<.001$).

4.3.1 Experiment 3a additional analysis

Despite overall performance in this experiment remaining significantly above chance, here I have followed the procedure used for the scrambled faces and conducted the analyses again whilst excluding those whose performance was below chance. In experiment 3a this resulted in the exclusion of 14 participants (of these 9 were recruited through the university and 5 through prolific) leaving a sample size of 130. Performance for all stimulus conditions were already significantly above chance in the original analysis and as such this is not reported again here.

A 2x2 within-subjects ANOVA using the factors *Face Type* (normal contour, blurred contour) and *Orientation* (upright, inverted) revealed a significant main effect of Orientation, $F(1, 129)=44.87$, $p<.001$, $\eta^2_p=.25$, indicating an overall robust inversion effect. A significant main effect of Face Type, $F(1, 129)=3.81$, $p=.049$, $\eta^2_p=.029$, was also found showing that overall recognition performance was higher for faces with normal contour ($M = .55$, $SD=.44$) than for faces with blurred contour ($M=.45$, $SD=.42$). No significant interaction was found, $F(1, 129)=1.28$, $p=.19$, $\eta^2_p=.012$ (see Figure 10).

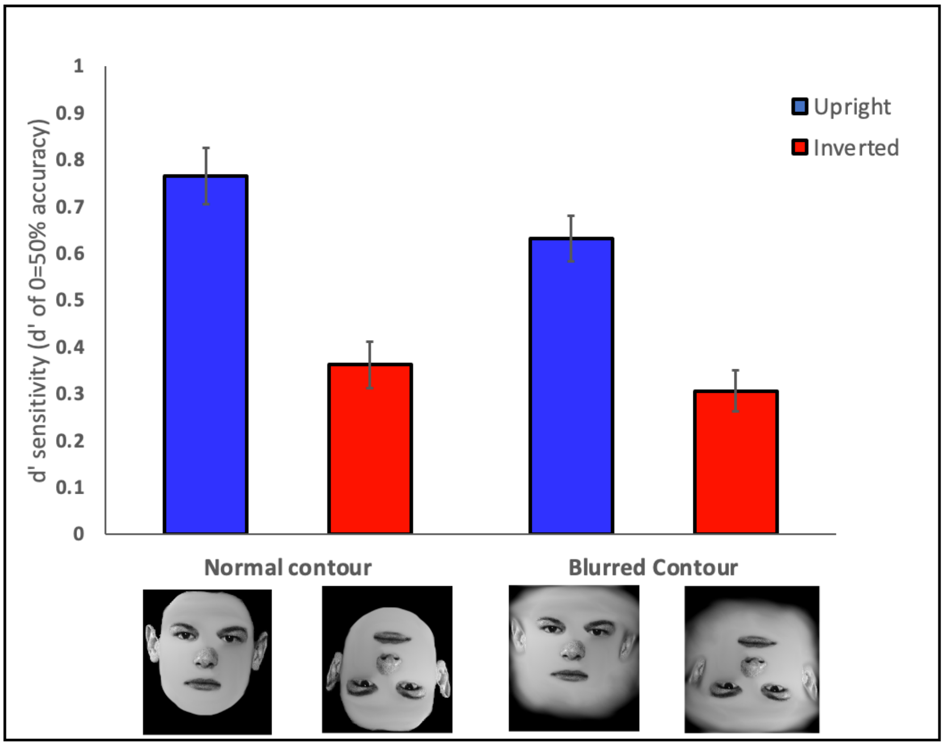

Figure 10. Graph reports the results from Experiment 3a (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, faces with either normal or blurred contour. The y-axis shows d′. Error bars represent s.e.m.

4.4 Discussion

In experiment 3a I have offered evidence to suggest that disruption of face contour through the blurring manipulation does not significantly affect the inversion effect. There is however a significant effect on overall performance as a result of the contour manipulation with performance being significantly reduced when the outline is altered. This suggests that while not specifically affecting the face inversion effect, the face contour may still play a role in overall face recognition performance for normal faces. However, one may argue that it is possible that the blurring manipulation disrupts not only the outline but also degrades the entire image due to the blending. Additionally, the face contour also contains spatial information which is part of what makes the faces distinct from one another and eliminating this may have made recognition more difficult beyond the alteration of the contour information. Another possible impact of blurring the outline is that, in addition to making the images less face-like, it also makes the images less like any recognisable stimulus which removes any sense of "objecthood" the image has. These considerations are addressed in experiment 3b for which I designed a new manipulation that aimed to maintain the integrity of the face stimuli and also preserve the information provided by the outline that makes the faces distinct from one another (but still eliminated the characteristic face outline). Experiment 3b was therefore a replication of experiment 3a in terms of behavioural paradigm but this time with a contour manipulation that would replace the original face outline with a novel face outline rather than blurring it in an attempt to remove it entirely.

4.5 Methods experiment 3b

4.5.1 Participants

In total, 144 participants (female=72, mean age=26.1, age range=18-71) took part in experiment 3b. In this case all participants were recruited through Prolific and were compensated for their participation. The sample size was determined from previous studies utilising the same face stimuli, counterbalancing of the participant conditions and stimuli, and behavioural paradigm (Civile et al., 2014; Civile et al., 2016).

4.5.2 Materials

Experiment 3b used two sets of stimuli. One set was the same normal-contour faces used in experiment 3a. The other set consisted of these original faces with an additional contour manipulation. For each face we generated a new spiked outline based on the original contour information. The faces had 8 spikes placed on the outside of the face spaced evenly around the outline with the existing spatial relationships between the features of face and the outline used to determine the size of the spike. The distance from the centre of the nose to the original outline was measured and this measurement was as the length of a given spike. These measurements always began at the centre of the nose and ended at the point on the outline where the spike was positioned. In this way I was able to create a unique new outline for every face and the variance in outlines between faces was determined by the variance that already existed in the normal-contour faces. Therefore, the information provided by the outline that helps distinguish one face from another is preserved whilst still removing the characteristic face outline that has been theorised to elicit holistic processing. The point at which the spikes returned to the face and joined one another was positioned slightly outside of the original outline in order to prevent an effect akin to the Kanizsa triangle illusion (Kanizsa, 1955) which results in a phenomenon whereby the face outline is seemingly generated by the eye as it automatically fills in the space between the bottom points of the spikes. With the new outlines in place, they were then smoothed to blend in with the existing face (in terms of both colour and texture) to create a cohesive face image. In the same manner as experiment 3a, the images were standardised to greyscale on a black background (see Figure 11). As before stimuli were counterbalanced across participants groups meaning each face identity was shown to every participant, but the orientation and contour were dependent on the participant group. The experiment was programmed and run on Gorilla.sc and participants used their own laptop or desktop devices.

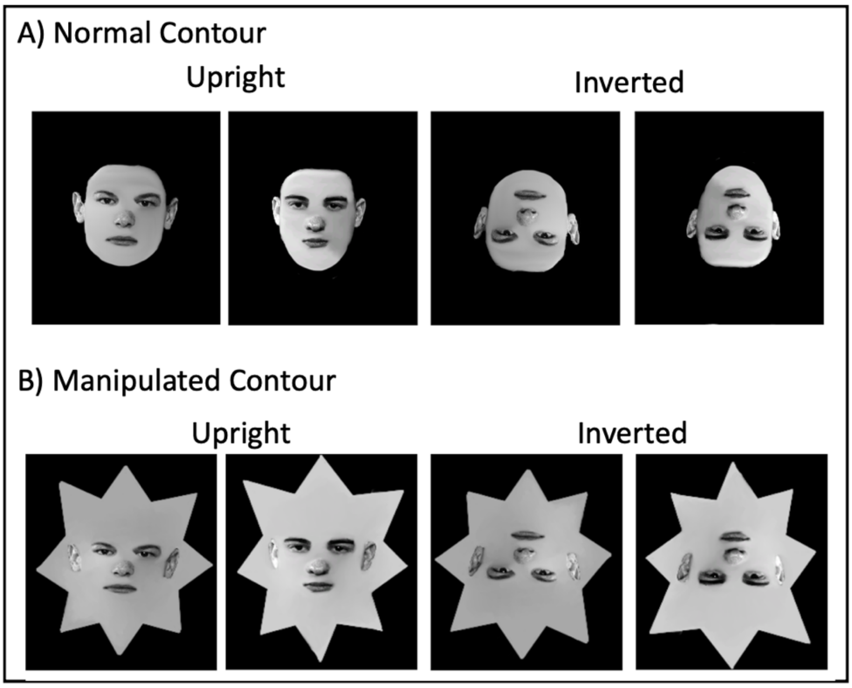

Figure 11. Examples of the stimuli used in Experiment 3b. Panel A shows faces with normal contour and panel B shows faces with manipulated contour.

4.5.3 Procedure

As in experiments 2 and 3a the behavioural task used was an old/new recognition task consisting of a study phase and a recognition phase. The study phase was run over 64 trials, each one began a 1s fixation cue in the centre of the screen followed by a face image presented for 3s. In line with the previous experiment, the faces were split evenly between normal-contour and new-contour faces with half of each set presented upright and half inverted and, these were presented intermixed, in random order. No response was required from participants during the study phase, and they were asked to memorise as many of the faces as possible. The recognition phase consisted of 128 trials with 50% of those involving the stimuli from the study phase and 50% involving novel stimuli presented one at a time at random order. Each trial began with a 1s fixation cue in the centre of the screen, followed by face stimulus shown for a maximum of 3s. Participants responded using the “X” and “.” keys to indicate whether or not they thought a given stimulus had been shown in the study phase (the meaning assigned to the keys was counterbalanced across participant groups). If no response was given after 3s participants were timed out and the next trial began automatically.

4.6 Results

A 2x2 within-subjects ANOVA using the factors *Face Type* (normal-contour, new-contour) and *Orientation* (upright, inverted) revealed similar results to that in 3a. A significant main effect of *Orientation*, *F*(1, 143)=42.00,p<.001, $\eta^2 p$=0.23, demonstrating the well-established face inversion effect. A significant main effect of *Face Type*, *F*(1, 143)=7.45, p=0.006, $\eta^2 p$=.05, was also found showing that overall recognition performance for normal-contour faces (M=0.56, SD=0.48) was significantly better than that for new-contour (M=0.40, SD=0.40). Crucially, no significant interaction (*Face Type x Orientation)* was found, *F*(1, 143)=2.57, p=0.11, $\eta^2 p$=0.02, indicating that both sets of faces showed a robust inversion effect, even though as in the case of Experiment 3a the inversion effect for normal faces with new-contour (M = 0.28, SD = 0.74) was numerically smaller than the one found for normal faces with normal contour (M = 0.42, SD = 0.92).

Importantly, comparison of stimuli against chance showed that performance for upright normal-contour faces (p<.001), upright new-contour faces (p<.001), inverted normal-contour faces (p<.001), and inverted new-contour faces (p<.001), were all significantly above chance. When all stimulus conditions were collapsed, overall performance was also found to be significantly above chance (p<.001).

4.6.1 Additional analysis

Again, additional analyses were performed excluding participants whose performance was below chance, for experiment 3b this meant 13 participants were excluded leaving a sample size of 131. Once again given that performance was already significantly above chance in the original analysis this comparison is not reported here.

A 2x2 within-subjects ANOVA using as factors *Face Type* (normal contour, new contour) and *Orientation* (upright, inverted) revealed a significant main effect of Orientation, *F*(1, 130)=42.56, p<.001, η^2_p=.24, again demonstrating the inversion effect. A significant main effect of Face Type, *F*(1, 130)=16.81, p<.001, η^2_p=.11, was also found, showing how overall recognition performance was higher for faces with the normal contour (M=.63, SD=.48) than for faces with the new contour (M=.47, SD=.42). In this case there was a significant interaction found, *F*(1, 130)=4.11, p=.045, η^2_p=.03, indicating differences in the inversion effect between the two face types. The breakdown of this interaction shows a large inversion effect for normal-contour faces with performance for upright faces (M=.86 SD=.72) significantly higher than that for inverted faces (M=.41 SD=.59), *t*(130)=5.75, p<.001, η^2_p=.18. For the new-contour faces there was a reduced but still significant inversion effect, again with performance for upright faces (M = .61 SD = .59) significantly higher than that inverted faces (M = .34 SD = .53), *t*(130)=4.08, p<.001, η^2_p=.10 (see Figure 12). Comparison of the upright stimuli showed that performance for upright normal-contour faces was significantly higher than that for new-contour faces, *t*(130)=4.12, p<.001, η^2_p=.106. No significant difference was found between the inverted stimuli of the two groups, *t*(130)=1.13, p=.25, η^2_p<.01.

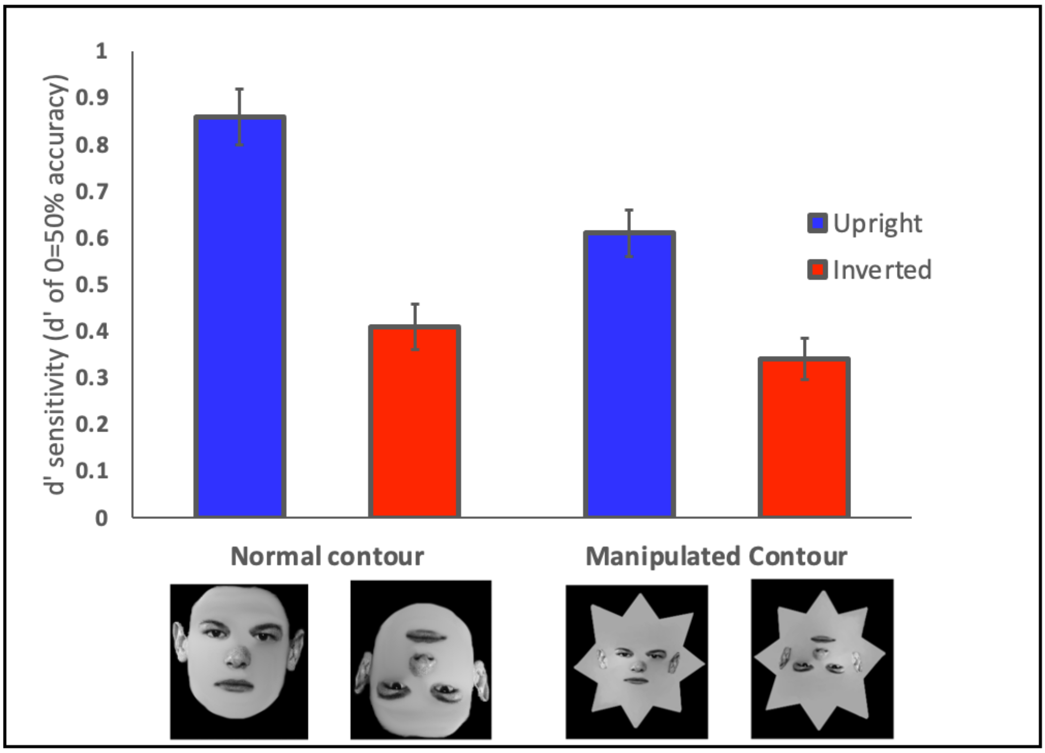

Figure 12. Graph reports the results from Experiment 3b (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, faces with either normal or manipulated contour. The y-axis shows d'. Error bars represent s.e.m.

4.7 Combined Original Analyses for Experiment 3a & 3b

Following the additional analyses of these two studies we are left with some slight disparity in the findings, in experiment 3a there was no significant interaction between Face Type and Orientation while a significant interaction is found for these factors in experiment 3b. It is possible that this is the result of the different contour manipulations given the potential issues that might exist in the stimuli which underwent the blurring manipulation; if the overall images in experiment 3a are degraded and have a reduced sense of objecthood then this might impair recognition performance for both the upright and inverted stimuli, overshadowing any orientation specific effect we might have expected to observe. However, it is potentially important to note that while the interaction seen in experiment 3b is significant, it narrowly meets the threshold for this at p=.045. This may indicate that the reduction in the inversion effect in this experiment is a less robust and reliable finding than that for the scrambled faces in experiment 2 and this may provide an alternative explanation for the disparity in results between experiments 3a and 3b. In order to gain insight into which of these explanations may be the most likely I conducted some additional analyses in which the data from experiments 3a and 3b are pooled to allow comparison of the manipulations and offer greater statistical power by way of a greatly increased sample size. Given the same behavioural paradigms (with the exception of the contour manipulation used) in experiments 3a and 3b as well as the matching participants numbers and similar results, I was able to pool the data across experiments to compare the contour manipulations and their impact on the inversion effect and overall face recognition performance.

A 2x2x2 mixed model ANOVA with the within-subjects factors *Face Type* (normal-contour, blurred/new-contour) and *Orientation* (upright, inverted), and the between-subjects factor *experiment* (experiment 3a, experiment 3b) revealed no significant main effect of *Experiment*, $F(1, 286)=0.803$, $p=0.371$, $\eta^2p<.01$. It also revealed no significant interaction between *experiment x Face Type*, $F(1, 286)=0.009$, $p=0.924$, $\eta^2p<.01$, or for *experiment x Orientation*, $F(1, 286)=0.037$, $p=0.848$, $\eta^2p<.01$. There was also no significant three-way interaction found, $F(1, 286)=0.042$, $p=0.683$, $\eta^2p<.01$. This series of results therefore suggest no evidence that the blurred-contour vs new-contour manipulation would impact differently on either overall performance or the inversion effect.

Importantly, the ANOVA revealed a significant main effect of *Orientation* (performance for upright better than inverted), $F(1, 286)=92.65$, $p<.001$, $\eta^2p=0.245$, and a significant main effect of *Face Type*, $F(1, 286)=12.85$, $p<.001$, $\eta^2p=0.043$ (higher overall performance for normal-contour faces (M=0.52, SD=0.51) vs blurred/new-contour faces (M = 0.41, SD = 0.44)). Critically, a significant interaction was found between *Face Type x Orientation*, *F*(1, 286)=3.95, p=0.047, $\eta^2 p$=0.014. To breakdown the *Face Type x Orientation* interaction paired t-tests were conducted which revealed a large inversion effect for normal-contour faces (M=0.40, SD=0.83), *t*(287)=5.50,p<.001, $\eta^2 p$=0.192 and a significantly reduced inversion effect for the blurred/new-contour faces (M=0.28, SD=0.74), *t*(287)=4.52, p<.001, $\eta^2 p$=0.130. In line with previous research (Civile et al., 2014a, 2016a), I compared performance for the upright and inverted stimuli individually between the two face types to investigate which stimuli are mainly driving the reduction of the inversion effect. Independent t-tests revealed that performance for upright normal-contour faces (M=0.73, SD=0.71) was significantly higher than that for upright blurred/new-contour faces (M=0.56, SD=0.59), *t*(287)=14.17, p<.001, $\eta^2 p$=0.04. No significant difference was found between inverted normal-contour faces (M=0.32, SD=0.59) and blurred/new-contour faces (M=0.27, SD=0.55), *t*(287)=1.54, p=0.215, $\eta^2 p$< .01. Indicating that as has been demonstrated in previous studies (and predicted by the perceptual learning account), the reduction in the inversion effects is driven mainly by differences in performance for upright faces.

4.7.1 Combined Additional Analyses for Experiment 3a & 3b

For the sake of completeness, I have further included a combined analysis of experiments 3a and 3b in which the participants whose performance was below chance were excluded. These results fall in line with the original combined analysis ad there was found a significant main effect of *Orientation*, *F*(1, 259)=85.87, p<.001, η^2_p=.25, demonstrating the inversion effect, and a significant main effect of *Face Type*, *F*(1, 259)=16.63, p<.001, η^2_p=.06 with overall recognition performance higher for faces with normal contour (M=.59, SD=.46) vs faces with blurred/new contour

(M=.46, SD=.42). A significant interaction was found between, *Face Type* and *Orientation*, *F*(1, 259)=5.41, p=.021, η^2_p=.021. Paired-samples *t*-tests were conducted to explore this interaction and revealed a large inversion effect for normal-contour faces (M=.42, SD=.82), *t*(259)=8.33, p<.001, η^2_p=.21 and a significantly reduced (but not eliminated) inversion effect for the blurred/new-contour normal faces (M=.28, SD=.75), *t*(259)=6.07, p<.001, η^2_p=.124 (see Figure 13).

To confirm the previous findings showing disruption to the upright stimuli is driving this reduction in the inversion effect we directly compared performance for the upright and inverted stimuli separately. Performance for upright faces with normal contour (M=.80, SD=.67) was significantly higher than that for upright faces with blurred/new contour (M=.60, SD=.58), *t*(259)=4.38, p<.001, η^2_p=.06. No significant difference was found between inverted faces with a normal contour (M=.38, SD=.56) vs blurred/new contour (M=.32, SD=.54), *t*(259)=1.38, p=.168, η^2_p<.01. Overall, these results are in line with those found in the original analyses.

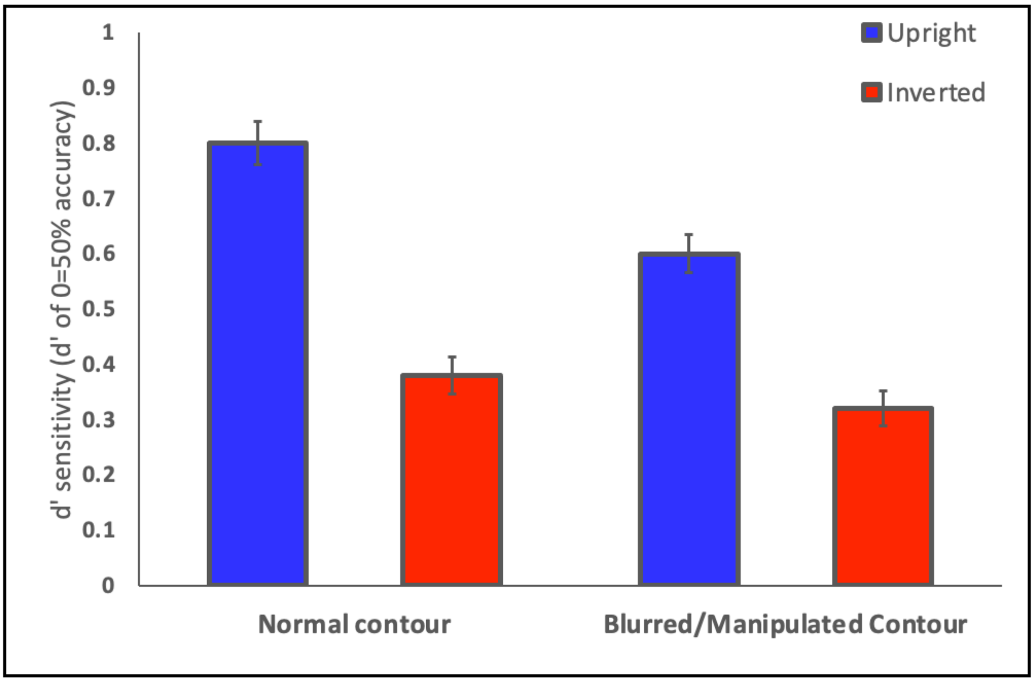

Figure 13. Graph reports the results from Experiments 3a and 3b combined (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, faces with either normal or blurred/manipulated contour. The y-axis shows d′. Error bars represent s.e.m.

4.7.2 Bayes Factor Analysis

Once again, I have conducted Bayes factor analysis, in this case on the difference between the d' values for upright and inverted faces (the inversion effect) comparing normal faces with normal contour vs normal faces with blurred/new contour for the additional analysis where Experiment 3a and 3b data were pooled together. The procedure used was that defined by Dienes (2011), using as a prior the interaction from Experiment 2, setting the standard deviation of p (population value | theory) to the mean of the differences in the inversion effect between normal vs blurred contour scrambled faces (0.19). We used the standard error (0.05) and mean difference (0.14) between normal contour and blurred/new contour normal faces from Experiment 3a and 3b combined. We assumed a one-tailed distribution for our theory and a mean of 0. This produced a Bayes factor of 19.78 providing strong evidence (greater than 10, for the conventional cut-offs see Dienes 2011, Jeffreys 1961) that the results from experiments 3a and 3b combined are in agreement with those of Experiment 2, showing that the reduction of the inversion effect when the face contour is manipulated is a reliable finding. Additionally, I conducted a Bayes Factor analysis using as a prior the mean difference between upright scrambled faces (normal contour vs blurred contour) in Experiment 2 (0.15). We then used the standard error (0.05) and mean difference (0.20) between upright normal faces (normal contour vs blurred/new contour) in Experiment 3a and 3b together. This produced a Bayes factor of 845.15, which is strong evidence (greater than 10, for the conventional cut-offs see Dienes 2011, Jeffreys 1961) that these results are also in agreement with those of Experiment 2, indicating that when the face contour is manipulated recognition for upright faces (either scrambled or normal) is reduced.

4.8 Discussion

We now have some discrepancy between the results of experiments 3a and 3b, while the original analysis from experiment 3b is consistent with that from experiment 3a, indicating that altering the face outline either by blurring or recreating it impairs overall recognition performance but does not reduce the inversion effect, when participants whose performance was below chance were excluded, a significant interaction emerged in experiment 3b. This new finding instead indicates that the new contour significantly reduced the inversion effect compared to normal faces. This is further supported by the analyses conducted on the data pooled across experiments (including when participants performing below chance were excluded) which also indicates that the contour manipulation results in a significant reduction of the inversion effect. Additional evidence for this finding is shown in the bayes factor analysis, this demonstrated that the results of the combined analysis are in line with the reduction of the inversion effect seen in experiment 2 and that it is impaired recognition of the upright faces as a result of the contour manipulation which is driving this effect.

It is interesting to note that despite including the same number of participants as experiment 2, neither experiment 3a nor 3b display a significant 2-way interaction when looked at originally unless participant who performed below chance are excluded (in spite of performance for all stimulus conditions and overall performance being above change in the original analysis). One possible explanation for this may be that the configural information between the internal features of the face remains entirely unaltered in these stimuli and perceptual expertise with this configural information is one of the main factors influencing face recognition and this is much more difficult to disrupt with a contour manipulation alone. Assessing whether this is the case is somewhat complicated, however, comparison of the strength of the inversion effect for each of these altered contour stimulus groups (scrambled blurred, normal blurred, and normal manipulated) may help to indicate whether it is the additional configural manipulation in experiment 2 which drives this difference (this would be suggested by a significantly lower inversion effect for the scrambled faces) or if another factor is at play (this would be suggested if no significant difference is found) . These additional comparisons are presented in the following pilot data. What is clear from the evidence presented in this chapter is that manipulating the contour of the face has an impact on reducing the inversion effect and overall recognition performance.

When the evidence provided by the experiments in this chapter and that from previous literature that has focused on manipulating the information available to participants are taken altogether, we can characterise the inversion effects based on the presence or absence of the first-order configural information, the face contour information, and the single feature-orientation information (this is the information relating to orientation of individual facial features, separate from the overall orientation of the whole face). Manipulating this last information by inverting half of the main features of the face has been shown to reduce the inversion effect when first-order information is intact (Civile et al., 2016) and eliminate the inversion effect when first-order information is disrupted through scrambling (Civile et al., 2014). In this way it can be considered the inversion effect persists to some extent in experiments 3a and 3b due to the unaltered single feature-orientation information and first-order configural information. In the same vein, the robust inversion effect found for scrambled face stimuli in Civile et al. (2016) can be explained by the relatively unaltered single feature-orientation information and face contour. The New Thatcherised face stimuli used in Civile et al. (2016) for which one of the eyes, one of the ears, and either the nose or mouth inverted in an otherwise upright face, demonstrate a reduced yet still significant inversion effect. The New Thatcherised faces ensure that whether they are presented upright or inverted, 50% of the single-feature information is altered. There is a range of research indicating that featural information is sensitive to inversion and this reduction in the inversion effect supports this and the idea that single-feature information is at least contributory to the face inversion effect. That the inversion effect is still significant under this manipulation can be interpreted as a result of the first-order configural and face contour information remaining intact in New Thatcherised faces. This aspect is the focus of the next chapter which aims to assess whether controlling for single-feature information when the contour manipulation is applied significantly impacts the inversion effect.

4.9 Pilot data: Contour manipulated scrambled faces vs normal faces inversion effect

The results from experiment 3a and 3b show that despite having the same number of participants as experiment 2 and using the same behavioural paradigm, a significant two-way interaction only emerged when participants who performed below chance were excluded, or the data was pooled across experiments. One potential explanation for this is that the configural information between the internal features of the face remains entirely unaltered in the normal faces; perceptual expertise with this configural information is thought to be one of the main factors influencing face recognition and this is much more difficult to disrupt with a contour manipulation alone compared to with a scrambling manipulation. Evidence in favour of this explanation can be seen in the following analysis which compares the size of the inversion effects for the manipulated contour scrambled and normal faces, with the usual caveats about comparing across experiments. A significantly stronger inversion effect for the normal faces would indicate that perhaps configural information is the main driver of the inversion effect, while comparable inversion effect would indicate that face contour plays an equally important role. Comparisons were made between the blurred contour scrambled faces from experiment 2 and the blurred contour normal faces in experiment 3a, the manipulated contour faces in experiment 3b, and finally the blurred/manipulated contour faces combined.

A 2x2 mixed model ANOVA was conducted with the within-subjects factor *Orientation* (upright, inverted) and the between-subjects factor *Experiment* (2, 3a). This revealed a significant main effect of *Orientation*, *F*(1, 286)=19.821, p<.001, η^2_p=.065, demonstrating an overall inversion effect. A significant main effect was also found for *Experiment*, *F*(1, 286)=25.631, p<.001, η^2_p=.293, indicating that overall performance differed between experiments. There was a significant interaction found between *Orientation* and *Experiment*, *F*(1, 286)=4.531, p=.034, η^2_p=.016 indicating a difference in the inversion effects for blurred contour scrambled and blurred contour normal faces. An independent samples t-test revealed that this significant interaction is due to a significantly larger inversion effect for blurred contour normal faces (M= .29, SD=.74) than for blurred contour scrambled faces (M=.10, SD=.76) *t*(286)=2.12, p=.034, η^2_p=.016.

A 2x2 mixed model ANOVA was conducted with the within-subjects factor *Orientation* (upright, inverted) and the between-subjects factor *Experiment* (2, 3b). This revealed a significant main effect of *Orientation*, *F*(1, 286)=18.704, p<.001, η^2_p=.061, demonstrating an overall inversion effect. A significant main effect was also found for *Experiment*, *F*(1, 286)=36.359, p<.001, η^2_p=.113, indicating that overall performance differed between experiments. There was a significant interaction found between *Orientation* and *Experiment*, *F*(1, 286)=4.027, p=.046, η^2_p=.014 indicating a difference in the inversion effects for blurred contour scrambled and manipulated contour normal faces. An independent samples t-test revealed that this significant interaction is due to a significantly larger inversion effect for manipulated contour normal faces (M= .28, SD=.74) than for blurred contour scrambled faces (M=.10, SD=.76) *t*(286)=2.01, p=.046, η^2_p=.014.

A 2x2 mixed model ANOVA was conducted with the within-subjects factor *Orientation* (upright, inverted) and the between-subjects factor *Experiment* (2, 3a/b). This revealed a significant main effect of *Orientation*, *F*(1, 430)=25.953, p<.001, η^2_p=.057, demonstrating an overall inversion effect. A significant main effect was also found for *Experiment*, *F*(1, 430)=39.249, p<.001, η^2_p=.084, indicating that overall performance differed between experiments. There was a significant interaction found between *Orientation* and *Experiment*, *F*(1, 430)=5.761, p=.017, η^2_p=.013 indicating a difference in the inversion effects for blurred contour scrambled and blurred/manipulated contour normal faces combined. An independent samples t-test revealed that this significant interaction is due to a significantly larger inversion effect for blurred/manipulated contour normal faces (M= .29, SD=.74) than for blurred contour scrambled faces (M=.10, SD=.76) *t*(430)=2.40, p=.017, η^2_p=.013.

That the inversion effect is so consistently larger for manipulated contour normal faces than manipulated contour scrambled faces provides evidence that configural information may be an important factor driving the inversion effect while face contour plays a considerably lesser role. To assess this claim further, future research could directly compare the inversion effects for manipulated contour scrambled and manipulated contour normal faces in a within-subjects design of the kind used in experiments 2, 3a, 3b, and 4 which would provide more substantial evidence about the relative strengths of the inversion effects that remain after contour manipulation. Further avenues of future research relating to this can be found in the general discussion of chapter 7.

Chapter 5 Manipulating the Contour of New Thatcherised Faces

5.1 Introduction to the experiments

In addition to the scrambling manipulation described in previous chapters, Civile et al. (2014a) also investigated the role of single-feature orientation information on the inversion effect. This was done firstly in combination with the scrambling manipulation through the creation of novel stimuli which they called "50% Feature-Inverted and Scrambled Faces". In addition to the sequential scrambling of the features, half of the features (one eye, one ear, and either the nose or mouth) were inverted, consequently the configuration was such that 50% of the features were inverted no matter the overall orientation of the face. The specific alteration of the features was again derived from prototype-defined groups with stimuli counterbalanced across participants. In an old/new recognition paradigm with upright and inverted stimuli it was observed that while scrambled faces alone still showed a robust inversion effect, this new manipulation resulted in the elimination of the inversion effect with upright faces impaired to the extent that there was no longer a difference between upright and inverted stimuli. This, along with the work of McKone and Yovel (2009) demonstrating that featural information is sensitive to inversion demonstrates the importance of single-feature orientation in the inversion effect. Replication of these findings was obtained by Civile et al. (2016a), and in their subsequent experiments they aimed to investigate whether disruption to second-order configural information (through the scrambling manipulation) is required to obtain this effect, or whether disruption to single-feature orientation alone is sufficient. They began with faces based on the Thatcher illusion in which two images of Margaret Thatcher are presented to viewers in the inverted orientation, one is an unaltered poster her face while on the other the eyes and mouth are inverted relative to the rest of the face. In the inverted orientation viewers detect little difference between the images but when they are returned to the upright orientation the manipulation of the “Thatcherised” face becomes incredibly obvious. The explanation offered for this is that when the face is in the inverted orientation, configural processing is reduced and discrete processing is utilised to a greater extent, making it more difficult to detect the rotation of the eyes and mouth. Conversely, when the face is presented in the upright orientation, the usual configural processing is reinstated, causing the rotation of the eyes and mouth to become apparent. (Thompson, 1980; Lewis & Johnston, 1997; Civile, McLaren, et al., 2016). This manipulation has the advantage of disrupting second-order configural information (due to the additional variance in these faces compared to a typical prototype) and single-feature orientation information while leaving first-order configural information relatively unaltered. However, the eyes have been shown to be highly salient features for face recognition (Ellis et al., 1979; Haig, 1984; Hosie et al., 1988) and as such inverting both eyes may be problematic beyond the expected disruption to second-order and featural information. To account for this, Civile et al. (2016a) modified the Thatcher manipulation in a similar manner to the stimuli in their 2014 study with one eye, one ear, and either the nose or mouth inverted in an otherwise upright face and dubbed these stimuli “New Thatcherised Faces”. In the old/new recognition paradigm the New Thatcherised faces displayed a significantly reduced inversion effect compared to normal faces but still a significant inversion effect remained. Taken together, the results from Civile et al. (2014a) and Civile et al. (2016a) indicate that both single-feature orientation information and first-order configural information play a role in the inversion effect.

In this chapter I explore further the impact of the contour manipulation on the inversion effect and overall recognition performance, and do so with the intention of investigating whether single-feature orientation plays a causal role in the inversion effect and whether disruption to this in combination with the face contour is able to influence the inversion effect, even in the context of unaltered first-order configural information. For this purpose, I revisited the New Thatcherised face stimuli used in Civile et al. (2016), applying to them the spiked contour manipulation from experiment 3b and using the old/new recognition task.

5.2 Methods experiment 4

5.2.1 Participants

In total, 144 participants (female=70, mean age=30.4, age range=18-64) took part in this experiment. Participants were recruited through Prolific and were compensated for their participation. The sample size was determined from previous studies utilising the same face stimuli, counterbalancing of the participant conditions and stimuli, and behavioural paradigm (Civile et al., 2014; Civile et al., 2016).

5.2.2 Materials

This experiment also consisted of two set of face stimuli; in this case both featured some manipulation to the original image. One set was the New Thatcherised faces generated by Civile et al. (2016), these began with 128 face images that were standardized to greyscale on a black background. The original face images were selected from the Psychological Image Collection at Stirling open database. The hair and neck were cropped out leaving the face contour and ears intact. Faces were then manipulated by rotating half of the internal features (one eye (including eyebrow), one ear, and either the nose or mouth) 180°. Four prototype categories of these faces were generated by selecting four different combinations of features to be rotated, in each case 50% of the features were inverted and all exemplars from a given category shared the same orientation of the features. As with the scrambled faces from experiment 2, these were then counterbalanced across participants groups to ensure that their findings were not the result of a given configuration being particularly easy or difficult for participants to recognise in the upright orientation. This completed the manipulation for the first set of stimuli. The other set began with these New Thatcherised faces and altered them further by applying the same spiked contour manipulation shown in experiment 3b (see Figure 14). This was again done using the spatial relationship between the features and the existing outline to determine the length of a given spike so as to preserve the information provided by the outline that makes the faces distinct from one another (see Figure 14). As before stimuli were counterbalanced across participants groups meaning each face identity was shown to every participant, but in this case the four prototype categories were also counterbalanced across participant groups. The experiment was programmed and run on Gorilla.sc and participants used their own laptop or desktop devices.

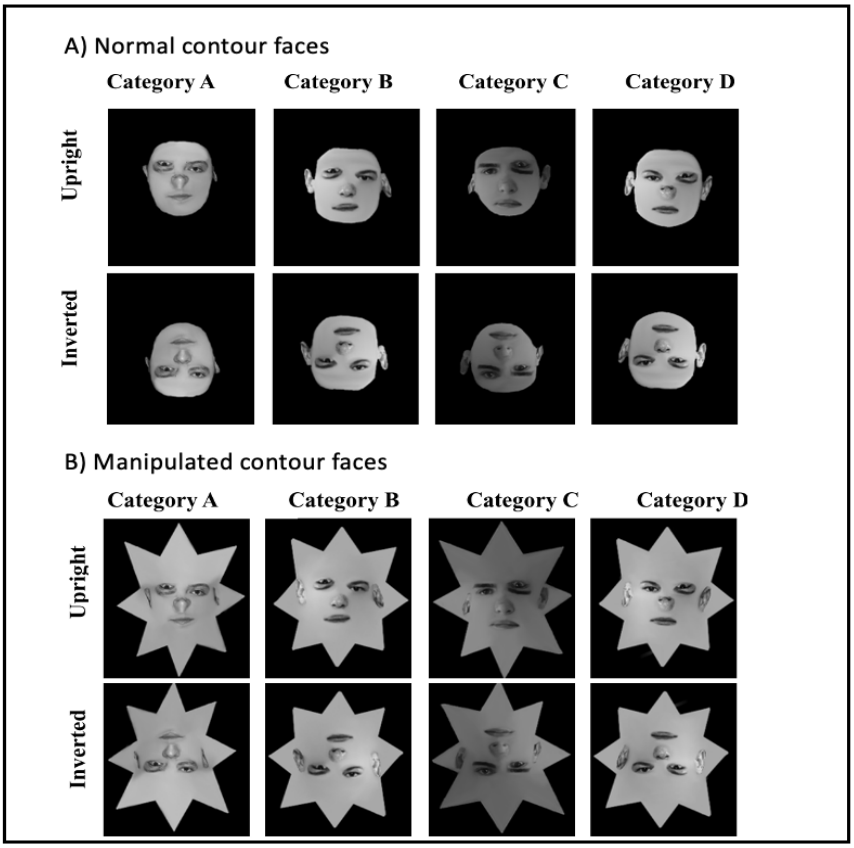

Figure 14. Examples of the stimuli used in Experiment 4. Categories specify the configuration of the internal features of the face. Panel A shows faces with normal contour and panel B shows faces with manipulated contour.

5.2.3 Procedure

In line with experiments 3a and 3b the behavioural task used was an old/new recognition task consisting of a study phase and a recognition phase. The study phase was run over 64 trials, each one began a 1s fixation cue in the centre of the screen followed by a New Thatcherised face image presented for 3s. As before, the faces were split evenly between normal-contour and new-contour faces with half of each set presented upright and half inverted and, these were presented intermixed, in random order. No response was required from participants during the study phase, and they were asked to memorise as many of the faces as possible. The recognition phase consisted of 128 trials with 50% of those involving the stimuli from the study phase and 50% involving novel stimuli presented one at a time in random order. Each trial began with a 1s fixation cue in the centre of the screen, followed by face stimulus shown for a maximum of 3s. Participants responded using the "X" and "." keys to indicate whether or not they thought a given stimulus had been shown in the study phase (the meaning assigned to the keys was counterbalanced across participant groups). If no response was given after 3s participants were timed out and the next trial began automatically.

5.3 Results

A 2x2 within-subjects ANOVA using as factors *Face Type* (normal-contour, new-contour) and *Orientation* (upright, inverted) revealed a significant main effect of *Orientation* $F(1, 143)=32.51$, $p<.001$, $\eta^2 p=0.185$, indicative of the inversion effect. There was also a significant main effect of *Face Type,* $F(1, 143)=3.98$, $p=0.047$, $\eta^2 p=0.027$, indicating that overall recognition performance for normal-contour New Thatcherised faces (M=0.25, SD=0.36) was better than that for new-contour New Thatcherised faces (M=0.17, SD=0.37). Critically, no significant interaction (*Face Type x Orientation*) was found, $F(1, 143)=0.17$, $p=0.68$, $\eta^2 p<.01$ indicating that a robust inversion effect was present irrespective of the contour manipulation. Notably, in this instance the inversion effect for the new-contour New Thatcherised faces was numerically larger (M=0.24, SD=0.68) than that for the normal-contour New Thatcherised faces (M=0.20, SD=0.68). Importantly, comparison of stimuli against chance showed that performance for upright normal-contour faces, $t(143)=7.82$, $p<.001$, $\eta^2 p=0.30$, upright new-contour faces, $t(143)=6.92$, $p<.001$, $\eta^2 p=0.25$, and inverted normal-contour faces, $t(143)=3.97$, $p<.001$, $\eta^2 p=0.09$, were all significantly above chance. However, performance for inverted new-contour faces, $t(143)=1.01$, $p=0.31$, $\eta^2 p<.01$, was not significantly above chance. Overall performance across all the four conditions was significantly above chance, $t(143)=9.53$, $p<.001$, $\eta^2 p=0.38$.

5.3.1 Additional analysis

In this case too I have performed some additional analyses in which participants whose performance was below chance were excluded, for experiment 4 this resulted in 27 participants being excluded leaving a sample of size of 117. A 2x2 within-subjects ANOVA using as factors *Face Type* (normal contour, new-contour) and *Orientation* (upright, inverted) revealed a significant main effect of Orientation, *F*(1, 116)=27.13, p<.001, η^2_p=.19 (the inversion effect), and a significant main effect of Face Type, *F*(1, 116)=4.49, p=.036, η^2_p=.037, showing overall recognition performance was higher for New Thatcherized faces with normal contour (M=.34, SD=.32) than that for New Thatcherized faces with the new contour (M=.23, SD=.34) (see Figure 15). No significant interaction was found, *F*(1, 116)=.03, p=.85, η^2_p<.01. Overall, these results are also in line with those found in the original analyses.

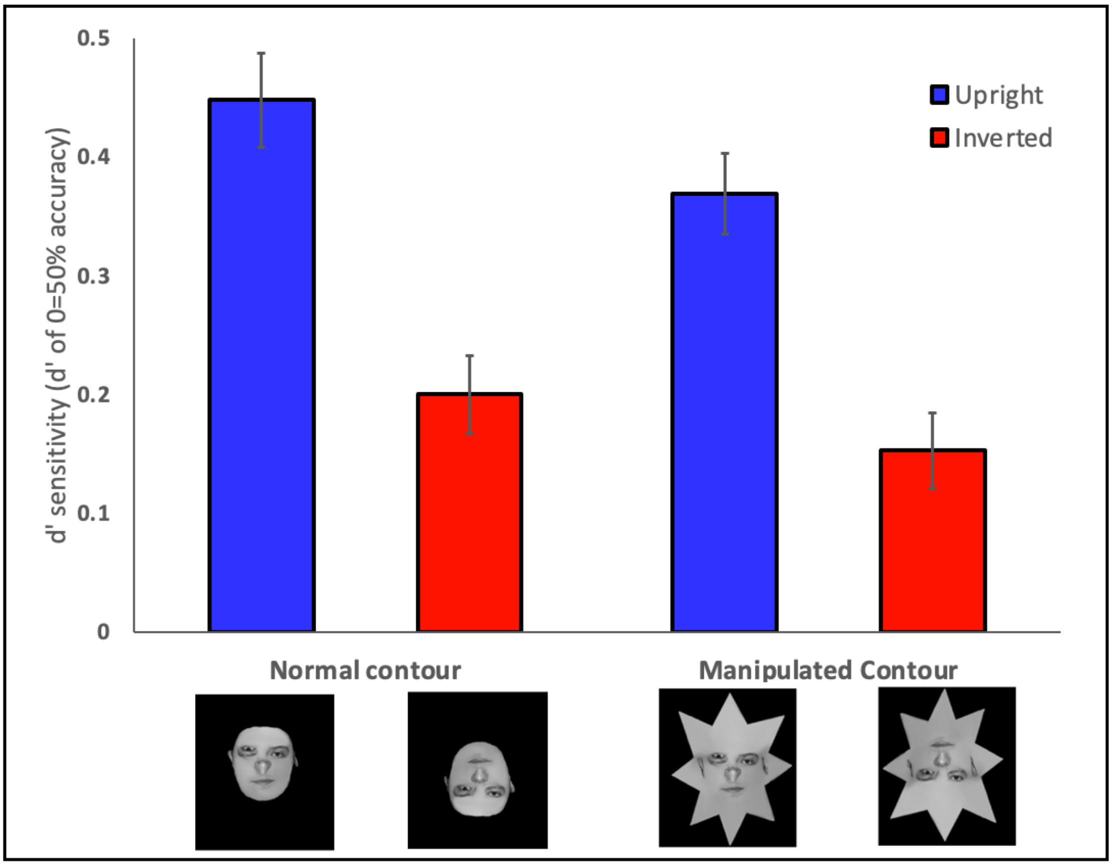

Figure 15. Graph reports the results from Experiment 4 (when participants whose performance was below chance were excluded). The x-axis shows the stimulus conditions, New Thatcherised faces with either normal or manipulated contour. The y-axis shows d′. Error bars represent s.e.m.

5.3.2 Bayes Factor Analysis

I again conducted a Bayes factor analysis, this time for the difference in overall performance between the New Thatcherized faces with normal contour and those with a new contour, using as a prior (population value | theory) the difference in overall recognition performance between normal and blurred/new contour faces in experiments 3a and 3b combined (0.13). I assumed a one-tailed distribution for our theory and a mean of 0. We then used the standard error (0.04) and mean difference (0.11) between overall performance for New Thatcherized faces with normal contour vs new contour. This produced a Bayes factor of 18.48 which provides strong evidence (greater than 10, for the conventional cut-offs see Dienes 2011, Jeffrey 1961) that these results are in agreement with those of Experiment 3a and 3b, showing that the reduction in overall recognition performance when the face contour is manipulated is a reliable finding.

5.4 Discussion

The findings outlined in experiment 4 provide evidence that for New Thatcherised faces, disruption to the face contour does not impact the inversion effect when single feature orientation information is controlled for (3 features are upright and 3 features are inverted) and the second-order configural information is disrupted due to the Thatcherisation manipulation. This finding is further established by the Bayes factor analysis which provided further evidence that the contour manipulation reduces overall performance for these stimuli. The experiment outlined in this chapter serves as an extension of the previous work on New Thatcherised faces conducted by Civile et al. (2016). Similar to the results reported in their study, here too we can observe a significant inversion effect in normal-contour New Thatcherised faces. In this instance it is suggested that the inversion effect found for the New Thatcherized faces in their studies could be explained by holistic information. In the theory laid forth by Hole et al. (1999), holistic information can be elicited by anything corresponding to the basic plan of a face, and it is holistic processing that establishes that it is a face that is being shown. This is part of the explanation given for their finding that photographic negatives show the composite face effect in much the same way as positives; negative faces still conform to this basic shape of a face and are therefore able to elicit holistic processing. Civile et al. (2016) add support for this in their proposal that the remaining inversion effect for New Thatcherised faces is the result of the relatively unaltered first-order configural information and the face contour evoking holistic processing in the upright faces which gives them an advantage over the inverted stimuli leading to a significant inversion effect. The role of first-order information in holistic processing has clear links to the idea that this is what denotes that it is a face that is being perceived; first-order information is common across all normal faces and it would therefore be logical to conclude that it is likely we use this information to identify faces as a stimulus class. The findings from experiment 4 that manipulating face contour and controlling single-feature information does not impact the inversion effect offer support for this interpretation that first-order information elicits holistic processing as this would explain the benefit shown for upright faces.

Despite not significantly reducing the inversion effect, this experiment does still provide some support for the idea that contour information plays an important role in face recognition as the manipulation did result in a reduction in overall performance. This result is also found in original analyses of experiments 3a and 3b, lending further support to this idea. There is however some tension between this finding and that for experiment 2 which did not show a reduction in overall performance for the blurred contour faces (although a one-tailed test would become close to significant). The initial hypothesis for this difference was that perhaps performance for the inverted faces in experiment 2 was low enough to supress any effect of the manipulation. With performance not significantly better than chance it seemed plausible that the stimuli had hit a floor effect meaning that any further reduction in performance for inverted stimuli as a result of the manipulation would be impossible. However, the additional analysis with participants whose performance was below chance excluded refuted this by demonstrating a significant reduction in the inversion effect, but no effect on overall performance, even when performance for the inverted stimuli had been brought significantly above chance level.

That a reduction in the inversion effect is observed only when below chance participants are removed is in line with the matching additional analyses for experiments 3a and 3b which also showed that when the poorly performing participants were removed, a significant reduction in the inversion effect was revealed. In all these experiments initial analysis indicated an effect on overall recognition performance rather than the inversion effect. It is therefore unclear from these experiments what the effect of contour manipulation is. This offers an avenue for future research which could directly investigate the effect of face contour manipulation on overall performance, the use of a delayed matching task here could allow much higher levels of overall performance and may therefore be a more appropriate measure for this purpose.

Chapter 6 Modulation of Perceptual Learning: Improving Face Recognition with Cathodal tDCS

6.1 Introduction to the experiments

It has been widely demonstrated throughout the tDCS literature that anodal stimulation delivered over the DLPFC can reduce the inversion effect through impairment for performance on upright faces. That anodal tDCS over the DLPFC can impair performance was initially observed by Ambrus et al. (2011) on categorisation learning of dot patterns using a prototype distortion task. This finding was not expected at the time given the previous literature which had largely found that anodal stimulation improved performance, while the opposite was true for cathodal stimulation (Nitsche & Paulus, 2000). However, Ambrus et al. showed that this is not always the case as in their experiment anodal tDCS delivered before and during the training phase resulted in decreased accuracy in the following categorisation task. A similar tDCS paradigm was subsequently adopted by Civile, Verbruggen et al. (2016) in their checkerboard categorisation task, as these stimuli too are derived from prototype-defined categories and required pre-exposure in a categorisation task to familiarise participants with them. Stimulation was applied during this task which was followed by an old/new recognition task investigating the checkerboard inversion effect. They too found that anodal stimulation impaired performance, in this case specifically for the upright checkerboards, resulting in a significantly reduced inversion effect. The explanation offered for this differential effect of tDCS on upright and inverted stimuli is explained in terms of impairment to perceptual learning as defined by the MKM model proposed by McLaren, Kaye and Mackintosh (1989).

The MKM model (more fully explained in Chapter 1) is derived from theories of associative learning and explains perceptual learning in terms of the differential latent inhibition of the common elements representing the stimuli. The model conceptualises stimuli as ‘units’ representing sets of features which operate in an error-correcting associative network in which the error term delta (∆) modulates the salience of the units and in doing so controls learning. Hence, the relative salience of a given unit depends on how well predicted it is by the other units active in the associative network; when it is well predicted the error term remains low and thus its salience is also low (latent inhibition), when it is not well predicted the error term is high and thus its salience is also high and it will readily form new associations (learning). Thus, the advantage conferred to upright exemplars from prototype-defined categories is explained in terms of the features common to many exemplars being well predicted with a low error terms and therefore low salience, while the features unique to an individual or few exemplars (which are the most useful in helping discriminate one exemplar from another) are not well predicted resulting in a high error term and thus, they are highly salient. This advantageous effect is however dependent on familiarity with the category in the upright orientation and does not persist when stimuli are inverted, resulting in the inversion effect. It is this effect that Civile, Verbruggen et al. (2016) argue is disrupted by anodal tDCS, increasing generalisation, and thus leading to the upright stimuli being robbed of the advantage usually bestowed upon them and reducing performance for these stimuli closer to the level typically seen for the inverted stimuli.

This effect has also been demonstrated extensively in the face recognition literature, as familiar prototype-defined stimuli, faces too are subject to the perceptual learning effect described in the MKM model and anodal tDCS has been demonstrated to impair performance for upright faces and reduce the inversion effect across a broad range of studies (Civile, McLaren et al., 2019; Civile, Waguri et al., 2020; Civile, Cooke et al., 2020; Civile, McLaren et al., 2020; Civile, Quaglia et al., 2021; Civile, McLaren et al., 2021; Civile & McLaren, 2022; Civile, McLaren et al., 2023). If anodal stimulation is able to reduce upright performance and the inversion effect the question remains whether it is possible to increase the inversion effect again using cathodal stimulation (this being the reverse polarity of anodal stimulation) and restore the perceptual learning effect for upright faces. This question is of particular interest in helping us to understand how tDCS is bringing about the effect that it does for face recognition. While I have focussed on the MKM based theory it is not the only one that has been proposed to explain the effects induced by anodal stimulation, there are several competing explanations of this tDCS induced effect. Some suggest that anodal stimulation alters the way in which participants view and perceive the presented stimuli. Scan path alteration is one example of this type of explanation, this theory states that anodal stimulation may cause participants to scan the faces differently to usual in the study phase and thus they pay minimal attention to salient features or miss details which may help them to distinguish faces from one another in the recognition phase. There is evidence to suggest that individuals each have a ‘preferred looking location’ for faces which in healthy participants is largely matched to their ‘optimal looking location’ and when forced to fixate on other areas of the face their recognition performance decreases (Peterson & Eckstein, 2013). It stands to reason therefore that if anodal tDCS alters a participant’s scan path this could result in the reduced accuracy we see for upright faces under tDCS. Alternatively, there are theories that suggest that perhaps the effects of anodal tDCS are the result of a reduced ability to encode in memory the stimuli that are presented in the study phase, thus impairing performance in the recognition phase. These types of account have a key aspect in common which is that they suggest that anodal tDCS results in a change of state that would prevent the faces in the study phase from being learned effectively and thus in the recognition phase performance is reduced. Theories such as these which predict that anodal stimulation will prevent faces from being encoded in memory or alter the way in which participants view and perceive the presented faces will also therefore predict that once anodal stimulation has been applied during the study phase, the participant will have lost the opportunity to learn the faces and thus remember them in the recognition phase, even if the effect of the stimulation is subsequently reversed.

In opposition to these theories, the perceptual learning-based account from the MKM model centres not on a change of state but rather on the participant's ability to utilise the information that they see and encode. This theory states that anodal tDCS can influence the error-based modulation of salience previously outlined in the MKM model. This means that while the faces shown in the study phase are perceived and encoded normally, when making judgements on whether the faces in the recognition phase are old or new, participants are unable to use the usually advantageous prediction error component to aid their decisions and make comparisons between the faces represented in memory and the faces presented in the recognition phase. As a result of this, unlike other theories, the model's explanation of the effect of anodal tDCS leaves room for the possibility that the effect can be reversed after stimulation has been applied; if anodal tDCS is in essence 'switching off' the error based modulation of salience and this is the cause of the impaired performance for upright faces then it should be possible to remove this impairment when this component is 'switched on' again.

Whether the impairment to recognition performance under anodal stimulation can be reversed is an important consideration when attempting to distinguish between explanations of the effect. If the tDCS effect is due to a change of state involving visual perception or encoding, then it would not be expected that additional stimulation of the reverse polarity would have any effect because the viewing and encoding of the original stimuli has already occurred. On the other hand, if the effect is found to be reversable this would provide evidence in support of a change in the way participants use the information contained within the face. It is this that I aimed to investigate in the following experiments, I begin with a pilot study in which I use two different tDCS montages, the first involved delivery of the conventional anodal montage during the study phase of an old/new recognition task followed by sham stimulation in the recognition phase. The second involved the same anodal montage in the study phase but was followed by the opposite cathodal montage in the recognition phase to assess whether the effect of the anodal tDCS can be reversed by the cathodal stimulation. The second experiment replicates these pilot conditions with the addition of a montage in which sham stimulation is delivered in both phases of the old/new recognition task, thus allowing the inversion effects in the active conditions to be compared to a baseline.

6.2 Methods experiment 5a

6.2.1 Participants

Overall, 64 participants (female=43, mean age= 20 years, age range=18-27 years) took part in experiment 5a, they were undergraduate students at the University of Exeter and participated either for monetary compensation or course credit.

6.2.2 Materials

The faces used in this study were the same as those used in experiment 1 consisting of 256 face images standardised to greyscale on a black background, originally taken from the Psychological Image Collection at Stirling open database. The faces were cropped such that the hair and neck were removed, leaving a standardised oval shape (see Figure 16). These stimuli were counterbalanced such that each participant saw all of the faces but in a different orientation/condition depending on their group. The experiment was programmed and run on Superlab 4.0b on an iMac desktop computer with participants positioned around 70cm from the screen.

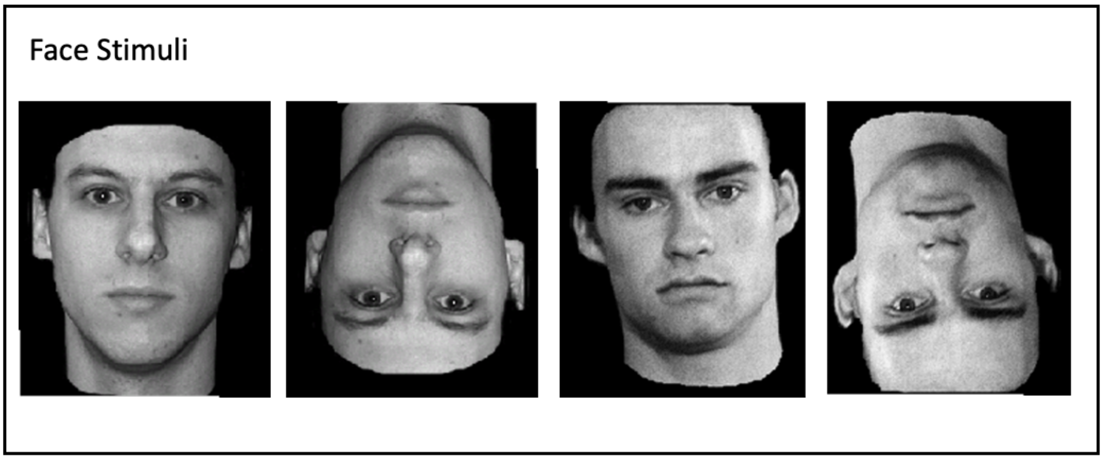

Figure 16. Examples of stimuli used in Experiment 5a and 5b. Upright and inverted faces with the hair cropped, standardised to greyscale on a black background

6.2.3 tDCS apparatus and montage

Participants were split into two conditions, one receiving anodal Fp3-Fp2 stimulation in the first half of the experiment and cathodal Fp3-Fp2 stimulation in the second half and the other receiving anodal Fp3-Fp2 stimulation in the first half and sham Fp3-Fp2 stimulation in the second half.

The tDCS system used was a Starstim tES-EEG system which delivered stimulation through two 35cm^2 electrodes encased in saline soaked sponges. These were positioned and held in place on the scalp through a cap with one at the target area of stimulation and another at the return location. The Starstim system allowed the use of a double-blind procedure in which a third-party experimenter (not actively running the study) loaded specific tDCS montages on the linked computer programme and provided a number for each participant which was used to determine whether they received anodal Fp3-Fp2, cathodal Fp3-Fp2, or sham Fp3-Fp2 stimulation. Participants experiencing the anodal Fp3-Fp2 and cathodal Fp3-Fp2 stimulation received 10 minutes of stimulation delivered at 1.5mA, while those experiencing sham received 30 seconds of 1.5mA stimulation followed by 0.1mA stimulation delivered for a total of 15 milliseconds spread over the 10-minute period (see Figure 17). Each condition began with a 5 second fade-in building intensity up to 1.5mA and ended with a 5 second fade-out reducing the intensity back down to 0.

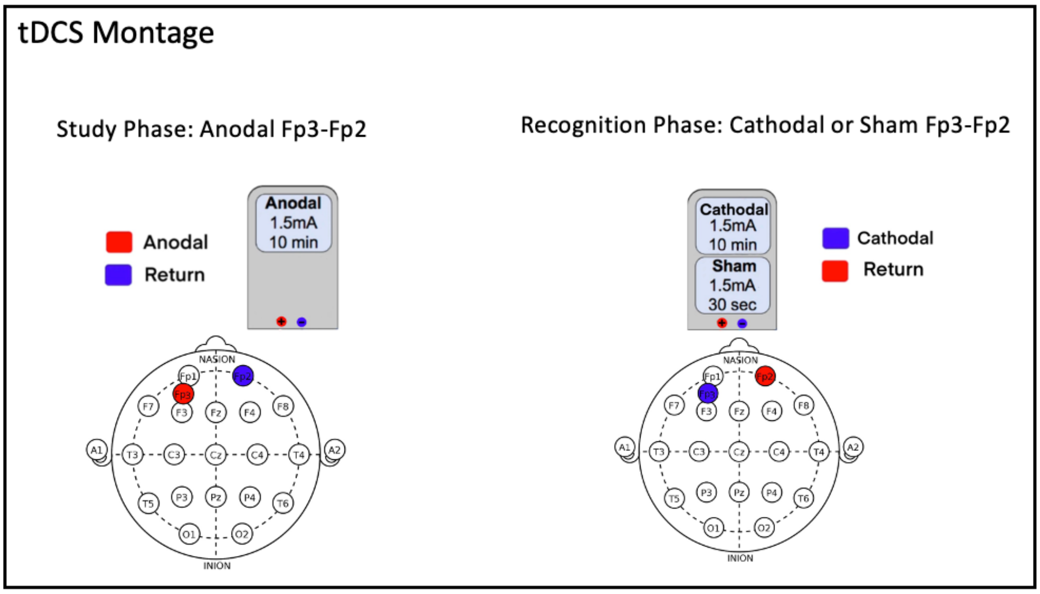

Figure 17. A representation of the tDCS montages used in Experiments 5a. In the study phase the tDCS montage used was always anodal Fp3-Fp2 and in the recognition phase the tDCS montage used was cathodal Fp3-Fp2 in the anodal-cathodal condition and sham Fp3-Fp2 in the anodal-sham condition

6.2.4 Procedure

The behavioural paradigm in this experiment consisted of an old/new recognition task with the first round of stimulation delivered during the study phase and the second delivered during the recognition phase. The study phase was run over 128 trials, each one began a 1s fixation cue in the centre of the screen followed by a face image presented for 3s. The faces were split evenly between upright and inverted and these were presented intermixed and in random order. No response was required from participants during the study phase, and they were asked to memorise as many of the faces as possible. Between phases participants had a 5-minute break while the second round of stimulation was initiated, and the next phase could begin. The recognition phase consisted of 256 trials, with 50% of those involving the stimuli from the study phase and 50% involving novel stimuli (also evenly split between upright and inverted orientations) presented one at a time at random order. Each trial began with a 1s fixation cue in the centre of the screen, followed by face stimulus shown for a maximum of 3s. Participants responded using the “X” and “.” keys to indicate whether or not they thought a given stimulus had been shown in the study phase (the meaning assigned to the keys was counterbalanced across participant groups). If no response was given after 3s participants were timed out and the next trial began automatically.

6.3 Results

A 2x2 mixed model ANOVA was conducted with the within-subject factor *Orientation* (upright, inverted) and the between-subjects factor *tDCS condition* (anodal-cathodal, anodal-sham). This revealed a significant main effect of *Orientation* $F(1, 62)=99.35$, $p<.001$, $\eta^2_p=.61$, demonstrating the inversion effect. There was no significant main effect of *tDCS condition* $F(1, 62)=1.07$, $p=.304$, $\eta^2_p=.017$ indicating that tDCS condition is not affecting overall performance. Importantly, there was a significant two-way interaction between *Orientation* and *tDCS condition* $F(1, 62)=4.17$, $p=.045$, $\eta^2_p=.063$.

As a follow up, paired samples t-tests were conducted on the inversion effect for each group. The anodal-sham condition showed a significant inversion effect with performance for upright faces higher than inverted M(difference)=.35, SD=.39, $t(31)=5.11$, $p<.001$, $\eta^2 p=.45$; as did the anodal-cathodal condition, although here the difference was larger, M(difference)=.53, SD=.32, *t*(31)=9.49, p<.001, $\eta^2 p$=.74 (see Figure 18). Comparison of these inversion effects in a between-subjects t-test revealed a significant difference between them with the inversion effect for the anodal-cathodal condition (M=.53, SE=.05) higher than for the anodal-sham condition (M=.35, SE=.06) *t*(62)=2.04, p=.045, $\eta^2 p$=.63. Additionally, performance for upright and inverted faces alone was compared across the tDCS groups based on the previous literature demonstrating that the tDCS procedure impacts upright but not inverted faces. Despite a numerical reduction in performance for upright faces in the anodal-sham condition (M=.55, SE=.06) compared to those in the anodal-cathodal condition (M=.70, SE=.06), there was no significant difference found (though there was a trend) *t*(62)=1.79, p=.07, $\eta^2 p$=.049. There was also no significant difference between the inverted faces in the anodal-sham condition (M=.19, SE=.05) compared to those in the anodal-cathodal condition (M=.16, SE=.04), *t*(62)=.424, p=.67, $\eta^2 p$=.049.

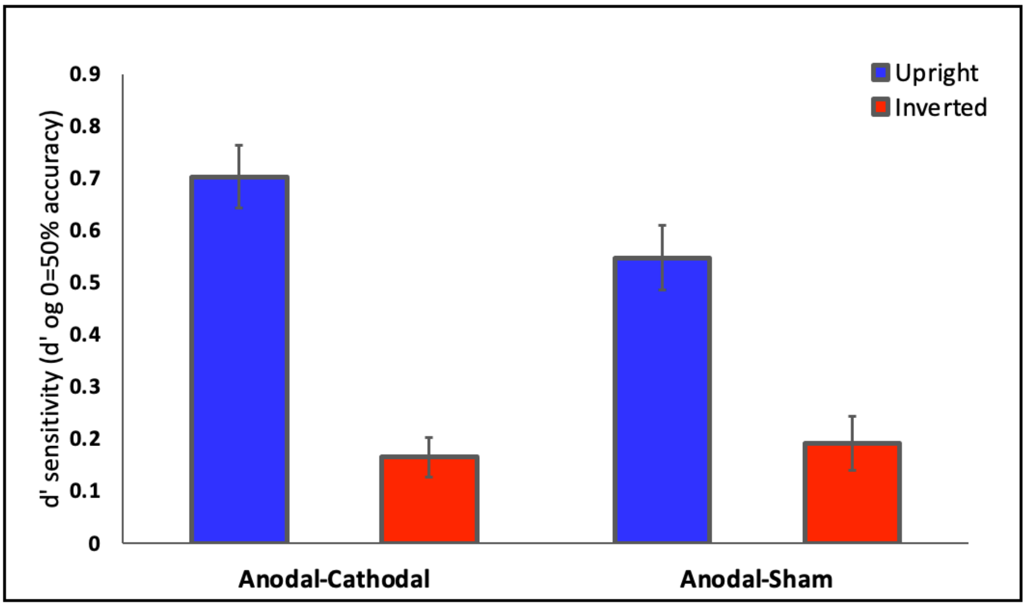

Figure 18: Graph reports the results from Experiment 5a. The x-axis shows the tDCS conditions, anodal-cathodal and anodal-sham. Both conditions have active stimulation during the first half. The y-axis shows d'. Error bars represent s.e.m.

6.4 Discussion

The results of this experiment are consistent with the idea that anodal Fp3-Fp2 tDCS stimulation significantly reduces the inversion effect for faces by impairing performance for upright faces. This stands in line with the extensive previous literature in this area. The novel finding from this experiment is that when anodal Fp3-Fp2 stimulation is followed by cathodal Fp3-Fp2 stimulation, there is a larger inversion effect compared to anodal Fp3-Fp2 stimulation alone and the upright faces are (numerically) driving this difference in inversion effects. This indicates that the tDCS induced effects for prototype-defined stimuli can be reversed and cathodal Fp3-Fp2 stimulation may be successful in restoring or even improving perceptual learning for faces. This finding offers preliminary evidence for an explanation of the tDCS effect in terms of increased generalisation due to modulation of perceptual learning, as described by the MKM model, rather than stimulation resulting in a change of state or encoding (i.e., change in scan path) which affects recognition performance. However, the extent to which the inversion effect was reduced by anodal Fp3-Fp2 stimulation and to which cathodal Fp3-Fp2 stimulation can reverse this effect cannot be accurately assessed in this experiment as we are unable to make a direct baseline comparison with the inversion effect for faces when participants have not been subject to any active stimulation (i.e., the typical face inversion effect). The following experiment extends this pilot experiment to include this condition with participants receiving sham stimulation for both tDCS montages (during the study phase and recognition phase). Thus, there are 3 tDCS procedures compared here, anodal-cathodal, anodal-sham, and sham-sham (the sham-sham group is also split with half using the Fp3-Fp2 set-up and half Fp2-Fp3 to maintain the double-blind procedure but with no expectation that there would be a difference between these).

6.5 Methods experiment 5b

6.5.1 Participants

Overall, 120 participants (female=86, mean age= 20.3 years, age range=18-46 years) took part in experiment 5a. As before, they were mainly students at the University of Exeter and participated either for monetary compensation or course credit.

6.5.2 Materials

The materials in this study were exactly the same as those used in experiment 5a.

6.5.3 tDCS apparatus and montage

Participants were split into three conditions: two were the same anodal-cathodal and anodal-sham conditions seen in experiment 5a, the additional condition was a sham-sham montage used to provide a baseline comparison. In this case the neuroConn DC-Stimulator Plus tDCS system was used which delivered stimulation through two 35cm^2 electrodes encased in saline soaked sponges. These were placed on the scalp with one at the target area of stimulation and another at the return location and held in place with adjustable head straps. The neuroConn system allowed the use of a double-blind procedure in which a third-party experimenter (not actively running the study) provided numerical codes which were used to determine whether they received anodal, cathodal, or sham stimulation. Participants experiencing the anodal and cathodal stimulation received 10 minutes of stimulation delivered at 1.5mA, while those experiencing sham received 30 seconds of 1.5mA stimulation followed by 0.1mA stimulation delivered for a total of 15 milliseconds spread over the 10-minute period. Each condition began with a 5 second fade-in building intensity up to

1.5mA and ended with a 5 second fade-out reducing the intensity back down to 0 (see Figure 19).

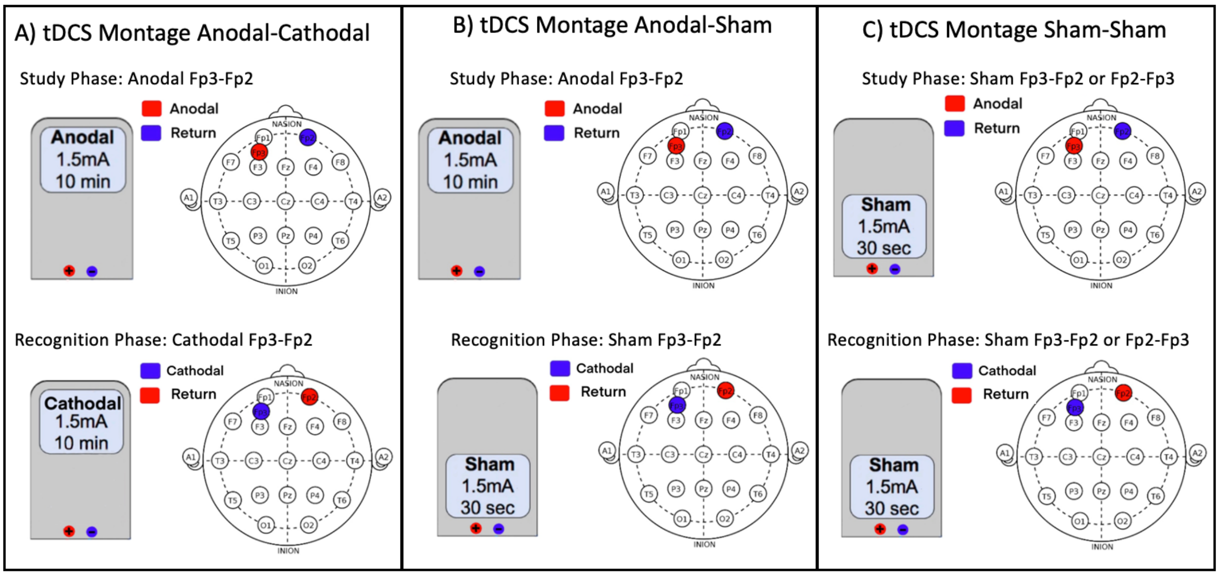

Figure 19. A representation of the tDCS montages used in Experiments 5b. Panel a. shows the set-up in the anodal-cathodal condition, in the study phase the tDCS montage used is anodal Fp3-Fp2 and in the recognition phase the tDCS montage used is cathodal Fp3-Fp2. Panel b. shows the set-up in the anodal-sham condition, in the study phase the tDCS montage used is anodal Fp3-Fp2 and in the recognition phase the tDCS montage used is sham Fp3-Fp2. Panel c. shows the set-up in the sham-sham condition, in the study phase the tDCS montage used with either sham Fp3-Fp2 or sham Fp2-Fp3 and in the recognition phase the tDCS montage used was also either sham Fp3-Fp2 or sham Fp2-Fp3.

6.5.4 Procedure

The behavioural paradigm used in this study was a direct replication of the old/new recognition task used in experiment 5a.

6.6 Results

A 2x3 mixed model ANOVA was conducted with the within-subject factor *Orientation* (upright, inverted) and the between-subjects factor *tDCS Condition* (anodal-cathodal, anodal-sham, sham-sham). This revealed a significant main effect of *Orientation F*(1, 117)=87.08, p<.001, η^2_p=.427, demonstrating an overall inversion effect. There was no significant main effect of *tDCS Condition F*(1, 117)=2.08, p=.129, η^2_p=.034, again indicating that tDCS does not have a blanket effect on overall performance. Importantly, there was a significant interaction between *Orientation* and *tDCS Condition F*(1, 117)=4.21, p=.017, η^2_p=.067.

Following this, paired t-tests were conducted on the inversion effect (upright compared to inverted faces) for each tDCS group. The sham-sham group displayed the expected large inversion effect with performance for upright faces higher than for inverted faces, M(difference)=.41, SD=.37, *t*(39)=7.03, p<.001, η^2p=.56. In the anodal-sham condition a reduced inversion effect was found, M(difference)=.22, SD=.44, *t*(39)=3.18, p=.002, η^2p=.21. In the anodal-cathodal condition there was an inversion effect of similar size (actually numerically larger) to the sham-sham condition, M(difference)=.51, SD=.52, *t*(39)=6.20, p<.001, η^2p=.49 (see Figure 20).

Independent t-tests were conducted to compare the size of the inversion effects in each tDCS group. There was a significant difference found in the inversion effect between the sham-sham condition and the anodal-sham condition, *t*(78)=2.10, p=.03, η^2p=.05, showing that in line with previous research, anodal tDCS reduces the inversion effect. There was also a significant difference between the anodal-sham and anodal-cathodal condition *t*(78)=2.64, p=.009, η^2p=.08, demonstrating that cathodal stimulation is able to counteract the reduction in the inversion effect that results from the anodal stimulation .There was no significant difference in the inversion effect between the sham-sham and anodal-cathodal conditions *t*(78)=.91, p=.36, η^2p=.01, indicating that cathodal stimulation can return performance back up to baseline (at least) following anodal stimulation.

Additionally, performance for upright faces alone was compared across the tDCS groups based on the previous literature demonstrating that the tDCS procedure impacts upright but not inverted faces. In the anodal-sham condition performance for upright faces (M=.45, SE=.05) was significantly reduced compared to the sham-sham condition (M=.66, SE=.06), *t*(78)2.43, p=.01, η^2p=.07, and the anodal-cathodal condition (M=.69, SE=.07), *t*(78)=2.57, p=.01, η^2p=.07. There was no significant difference between the inverted faces in the anodal-sham condition (M=.23, SE=.05) compared to the sham-sham condition (M=.25, SE=.05), *t*(78)=.351, p=.73, η^2p<.01, or the anodal-cathodal condition (M=.19, SE=.04), *t*(78)=.588, p=.56, η^2p<.01.

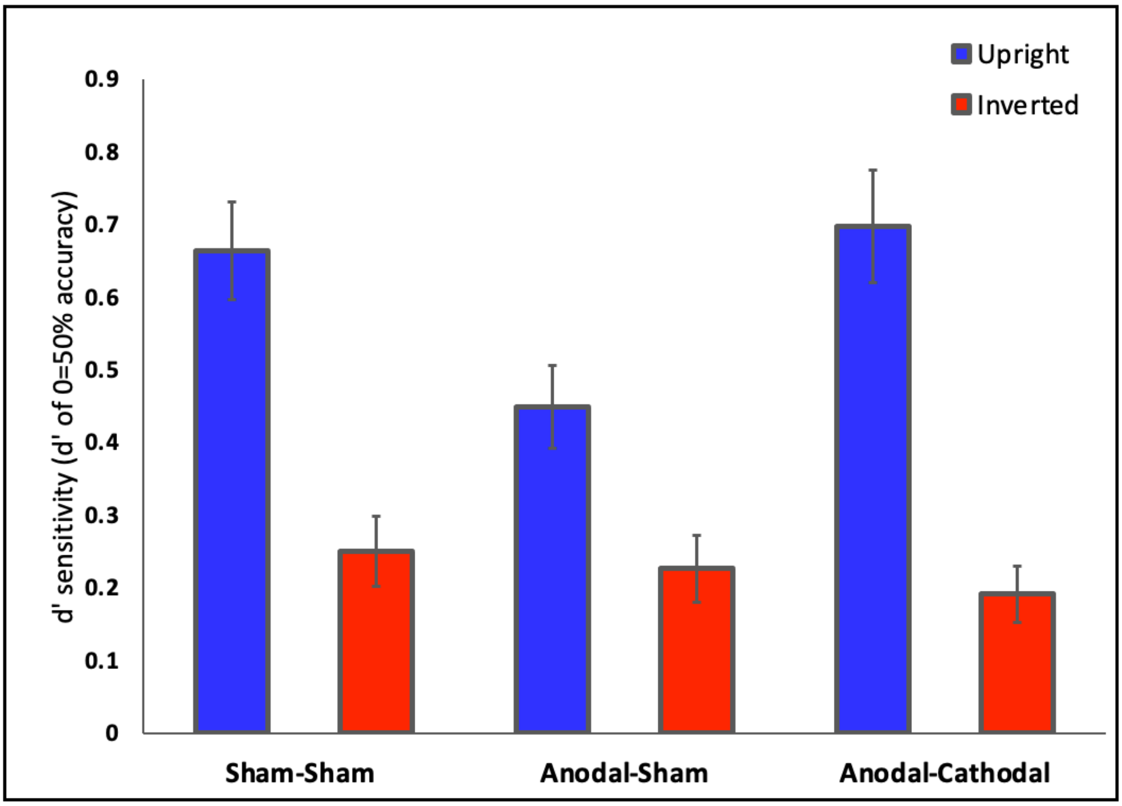

Figure 20: Graph reports the results from Experiment 5b. The x-axis shows the tDCS conditions, sham-sham, anodal-cathodal, and anodal-sham. The y-axis shows d'. Error bars represent s.e.m.

6.6.1 Combined Analysis

Given the equivalent behavioural paradigms used in experiments 5a and 5b as well as the matching participants numbers and similar results, I was able to pool the data across experiments to compare the results in the anodal-cathodal and anodal-sham conditions. Sham-sham is not included in the analysis as it was not a condition present in the pilot experiment 5a.

A 2x2x2 mixed model ANOVA was conducted with the within-subject factor *Orientation* (upright, inverted) and the between-subjects factors *tDCS Condition* (anodal-cathodal, anodal-sham) and *Experiment* (5a, 5b). This revealed a significant main effect of *Orientation* *F*(1, 140)=126.19, p<.001, η^2_p=.474, demonstrating an overall inversion effect. There was a near significant main effect of *tDCS Condition* *F*(1, 140)=3.828, p=.052, η^2_p=.027, but this is due to the effect on upright faces and no significant main effect of *Experiment* *F*(1, 140)=.063, p=.802, η^2_p<.001, indicating that the results were consistent across both experiments. There was a significant two-way interaction between *Orientation* and *tDCS Condition* *F*(1, 140)=10.472, p=.002, η^2_p=.070, showing that the tDCS differentially affects the upright and inverted faces. There was no significant two-way interaction between *Orientation and Experiment* *F*(1, 140)=1.315, p=.254, η2p<.01. There also no significant three-way interaction *Orientation x tDCS condition x Experiment* *F*(1, 140)=.488, p=.486, η2p<.01, indicating that the interaction between *Orientation* and *tDCS Condition* was consistent across experiments.

The significant interaction between orientation and tDCS condition was further investigated with a series of t-tests. Paired samples t-tests revealed that there was a significant inversion effect (with upright greater than inverted) in both the anodal-cathodal condition M(difference)=.52, SD=.44, *t*(71)=10.10, p<.001, $\eta^2 p$=.59 and anodal-sham condition M(difference)=.28, SD=.42, *t*(71)=5.64, p<.001, $\eta^2 p$=.31. Comparison of these inversion effects in an independent samples t-test revealed a significant difference between them, with the inversion effect in the anodal-cathodal condition (M=.52, SD=.44) significantly larger than that in the anodal-sham condition (M=.28, SD=.42) *t*(142)=3.35, p=.001, $\eta^2 p$=.14. As before this difference in the inversion effects was driven by reduced performance for the upright faces in the anodal-sham condition compared to the anodal-cathodal condition M(difference)=.21, SD=.52, *t*(71)=3.39, p=.001, $\eta^2 p$=.139 while there was no significant difference in the inverted stimuli for these groups M(difference)=-.03, SD=.32, *t*(71)=-.84, p=.402, η^2p=.009

6.6.2 Bayes Factor Analysis

Here too I provide Bayesian statistics to complete the comparison across experiments. According to the procedure devised by Dienes (2011) we conducted a Bayes analysis on the difference between inversion effects in anodal-sham and anodal-cathodal. We used as the *priors* the differences found in experiment 5a, setting the standard deviation of p (population value | theory) to the mean for the difference between the face inversion effect in anodal-sham group vs that in the anodal-cathodal group (0.18). We used the standard error (0.1) and mean difference (0.28) between the face inversion effect in the anodal-sham group vs. that in the anodal-cathodal group. This gave a Bayes factor of 19.31, which is very strong evidence (greater than 10, for the conventional cut-offs see Jeffrey, 1961; Dienes 2011) that these results are in line with the pilot.

We then conducted a Bayes analysis on the difference between upright faces in anodal-sham and anodal-cathodal. We used as the *priors* the differences found in experiment 5a, setting the standard deviation of p (population value | theory) to the mean for the difference between the upright faces in anodal-sham group vs that in the anodal-cathodal group (0.18). We used the standard error (0.1) and mean difference (0.28) between the upright faces in the anodal-sham group vs. that in the anodal-cathodal group. This gave a Bayes factor of 31.09, which is again very strong evidence that these results are in line with the pilot.

6.7 Discussion

These results firstly contribute to the existing body of work showing that anodal tDCS stimulation significantly reduces the inversion effect and does so through impairment to the upright stimuli. This finding can be explained in terms of the anodal Fp3-Fp2 stimulation disrupting the perceptual learning mechanism described in the MKM model and in doing so reducing the advantage usually seen for upright faces. In line with the pilot in experiment 5a these findings also indicate that cathodal Fp3-Fp2 stimulation is able to reverse the effects induced by the anodal stimulation and improve performance for upright faces such that it is no longer significantly different from baseline performance (seen in the sham-sham condition). This lends further support to the interpretation that tDCS modulates perceptual learning and now shows that it can do so in either direction depending on the polarity of the stimulation used. The Bayesian statistics conducted across experiments further confirm that there is a highly reliable difference between the reduced inversion effect in the anodal-sham condition and the inversion effect in the anodal-cathodal condition, solidifying these findings. As previously discussed, these results also support the interpretation that the tDCS effects result in changes in the type of information participants use to discriminate stimuli, increasing generalisation and in some sense making them more “alike” which make discrimination more difficult, particularly for the upright stimuli with which we would usually use our extensive expertise with certain information to aid recognition. This helps to refute the idea that anodal tDCS stimulation results in a change of state in participants which affects the way they encode the stimuli in the study phase because if this were the case then using a different form of encoding in the recognition phase due to cathodal stimulation should not improve recognition of faces which had already been encoded under anodal tDCS. In this way our results help further our understanding of the way in which tDCS stimulation impacts the inversion effect and face recognition.

Chapter 7: General Discussion

In this final chapter I will summarise the key findings of the experiments outlined in this thesis and consider their implications in terms of the central debates existing in the face recognition literature. These primarily relate to the specificity vs expertise accounts of the face inversion effect and the types of information that contribute to it. Some additional analyses and pilot work will be reported which I do not yet consider strong enough to present as stand-alone work but none-the-less offer interesting hints for future research.

7.1 Implications for the specificity vs expertise debate

The key debate in the face recognition literature since the initial discovery of the face inversion effect has been whether or not faces are special? More specifically, does processing for faces occur through some neural mechanism specific to faces (Yin, 1969), or is it the result of our life long expertise with them (Diamond & Carey, 1986)? That an inversion effect is reported for both faces and checkerboards in the sham condition in experiment 1 of chapter 2 contributes to the expansive body of work indicating the perceptual expertise gained through perceptual learning plays an important role in the inversion effect. This was initiated by McLaren (1997) and his work on prototype-defined checkerboards which demonstrated that artificial stimuli can produce an inversion effect provided that pre-exposure has occurred to allow for perceptual learning. This was further corroborated by the tDCS effects demonstrated in Civile, Verbruggen et al. (2016) showing that, as for faces, checkerboards too are impacted by the perceptual learning modulation that results from anodal tDCS. In addition to demonstrating that a shared perceptual learning component between faces and other prototype-defined familiar stimuli exists, the findings from experiment 1 also offer support for the possibility that there is still a role for facial specificity in the face inversion effect. In line with the findings from Civile, Quaglia, et al. (2021) there was found to be a significant difference between the entirely eliminated inversion effect for checkerboards and the reduced inversion effect for faces under anodal tDCS. Contrary to what has previously been suggested, the combined findings show strong evidence that this is not the result of checkerboard recognition being more difficult than faces; overall performance for face and checkerboards was not found to be significantly different in either study and comparison of the upright stimuli solidifies this position, again showing no significant difference. These results were show to be highly reliable through the Bayesian statistics using as priors Civile, Quaglia, et al.'s data and showing strong evidence of the effect. This could support the notion that in addition to perceptual learning there is an additional aspect to the inversion effect for faces which is unaltered by the tDCS, allowing us to integrate the expertise and specificity accounts which have long been in contention with one another.

However, the key factor that remains unaddressed by either experiment 1 or Civile, Quaglia, et al. (2012) is the vast difference in expertise that participants have for faces and checkerboards. Our expertise for faces is so great that most people do it entirely unconsciously in their everyday life without ever having to think about it, this lifelong experience can hardly be said to compare to the relatively short training session delivered to participants to make them familiar with checkerboards. Rather than any aspect of facial specificity, it is possible that tDCS is simply unable to affect the very high level of expertise that we have for faces to the same extent that it can affect the newly gained expertise for checkerboards. There are some important points to consider when exploring this explanation, including as mentioned above the fact that overall performance, and performance levels for the upright stimuli did not differ between groups, indicating a comparable level or recognition between them. Crucially though if we consider the explanation that the effect of tDCS is to disrupt perceptual learning we must acknowledge that participants have had far more opportunity for perceptual learning about faces than they have for checkerboards and this is not a factor that is easily ignored. Later in this discussion I will explore some future directions for research and there I suggest a study which would help to address whether this level of disparity in expertise may be contributory to the differential effects of anodal tDCS for faces and checkerboards. The basis for this study is increasing the amount of time that participants have to gain expertise for checkerboards by giving them more extensive training over days or weeks in a variety of learning tasks and investigating the impact this has on the inversion effect in a subsequent recognition task under anodal tDCS. An increase in the inversion effect that remains after anodal tDCS for checkerboards would indicate that it is the sheer level of expertise that underlies the difference between faces and checkerboards in experiment 1, while finding no difference would indicate support for a potentially face-specific component that is not able to be affected by the tDCS procedure.

7.2 Implications for studies of the face inversion effect

7.2.1 Specificity (Yin, 1969; Farah, 1995)

Unlike so many of the studies demonstrating the role of perceptual expertise in the inversion effect, I do not necessarily discredit the specificity account put forth by Yin (1969) on this basis. The evidence from experiment 1 does allow for the potential that there is a face-specific neural mechanism contributing to the inversion effect. It is clear from experiment 1 that there are differences in the extent to which anodal tDCS can affect the inversion effect for faces and other prototype-defined familiar stimuli and it is entirely plausible that this is reflective of a face specific mechanism. This would additionally support the findings relating to acquired prosopagnosia shown in Farah, et al. (1995) and perhaps the component of the inversion effect for face unaltered by the tDCS relates to an area of the brain that had been damaged in LH. While there is not specific evidence that whatever drives the remaining inversion effect for faces during tDCS is face-specific, neither do I find evidence to suggest the need to refute this explanation.

7.2.2 Expertise with prototype-defined stimuli (Diamond & Carey, 1986; Gauthier & Tarr, 1997)

Given that we have found evidence for a robust inversion effect (of the same level as that found for faces) for checkerboards stimuli (a familiar prototype-defined category) in the sham condition of experiment 1 we can offer support for Diamond and Carey's (1986) theory that the face inversion effect is related to expertise. This is further solidified by the evidence that modulation of perceptual learning through anodal tDCS is able to eliminate this inversion effect for checkerboards. The comparable inversion effects found for dogs with dog breeder participants and faces in their experiment matches well with our checkerboards and face inversion effect comparisons. This is also the case for the greebles generated by Gauthier and Tarr (1997), these too are artificially created, prototype define-stimuli and in that sense are analogous with checkerboards. Our findings also refute the criticisms that the greeble inversion effect may be due to their creature-like nature and similarity to faces in design, checkerboards do not share these qualities and yet we find supporting evidence with the none-the-less. In contrast to the interpretations of the results in these studies though I do not take evidence for the expertise account to necessarily refute the specificity account but rather offer an interpretation in which they can co-exist.

7.2.3 Expertise in terms of perceptual learning and the MKM model (McLaren, 1997; Civile, Zhao, et al., 2014; Civile, Verbruggen, et al., 2016; Civile, et al., 2018)

The results of experiment 1 continue to find evidence that perceptual learning, through exposure to prototype-defined stimuli, plays a role in the inversion effect. This is explained in the MKM model (McLaren, Kaye & Mackintosh, 1989) in terms of repeated exposure to a prototype-defined stimulus category reducing the salience of the common features and increasing the salience of the unique features, thereby aiding discrimination. This was demonstrated in McLaren's (1997) work with checkerboards, demonstrating that pre-exposure to these stimuli when they are derived from a prototype category results in better discrimination of them from other members of that category. Civile, Zhao, et al. (2014) extended this finding to a recognition paradigm a found similar results, that is that pre-exposure to a category of checkerboard in a categorisation task leads to better recognition performance for that category in old/new recognition task. As predicted by the MKM model in both these studies the advantage conferred by perceptual learning was only present for the upright exemplars (despite checkerboards being non-mono-oriented) and was lost when the checkerboards were inverted. The tDCS results presented by Civile, Verbruggen, et al. (2016) that anodal stimulation can eliminate the inversion effect for checkerboards by impairing performance for the exemplars provided the first clear evidence that tDCS can modulate perceptual learning in this way. The extension of this finding to faces in Civile, et al. (2018) confirmed this result, showing a significant reduction in the face inversion effect during anodal tDCS compared to sham. The results from experiment 1 further contribute to this body of work, providing support for the findings of this previous literature that both the face and checkerboard inversion effect is reduced by the tDCS procedure due to the modulation of perceptual learning for the upright stimuli.

7.3.1 The specific tDCS procedure, active control does not produce a reduction in the inversion effect (Civile, et al., 2018; Civile, McLaren, et al., 2021)

One of the factors explored in experiment 1 was the use of an active control condition to determine whether the effect of the tDCS stimulation is specific to the Fp3-Fp2 montage. This completed a series of active control used in perceptual learning studies, active control refers to stimulation of the same intensity and duration as the target stimulation but delivered at different channels. Civile, et al. (2018) used an rIFG-Fp1 active control in their study of the face inversion effect and discovered that not only did it not produce the reduction of the inversion effect seen for the Fp3-Fp2 montage but it also did not differ significantly from sham, indicating that the modulation of perceptual learning is not a general tDCS effect. Following this Civile, McLaren, et al. (2021) used a PO8-Fp2 in the investigation of the composite face effect and again found that it was not different from sham in its ability to modulate perceptual learning in this task, this allows us to more specifically say the effect is not solely to do with the return electrode placement at Fp2. This set-up over the occipital-temporal area is particularly important as there have been conflicting findings about the effect of delivery tDCS to this region. This area is often chosen based on the N170 ERP component, which has been linked to face-sensitivity. Yang et al. (2014) also looked at the composite face effect using electrodes at P7 and P8 (running P8-P7, P7-P8, and sham conditions) with active stimulation delivered a 1.5mA for 15 minutes. They utilised a within-subjects design in which participants engaged in the same task three times under the three different tDCS conditions (each separated by at least 3 days). The composite face effect was studied using a composite face task. Here faces are comprised of the top half of one face and the bottom half of another and they used a full design including congruent and incongruent trials and aligned and misaligned faces. Participants were presented two composite faces in succession and asked to judge whether the top half of the first face matched the top half of the second face. Congruent trials are characterised by the bottom half of the face supporting the same/different decision (i.e. where the top halves are the same the bottom halves are too and vice versa) while the opposite is true for incongruent trials. The typical composite face effect was found in the sham condition with accuracy higher for congruent than incongruent trials in the aligned condition and this difference reduced in the misaligned condition; subtracting the congruency effect in the misaligned and aligned trials offers an index of the composite face effect. In both active conditions this index was significantly reduced compared to sham, indicating that tDCS (regardless of polarity) is able to modulate the composite face effect when delivered over occipital-temporal areas. This finding however failed to be replicated in other similar experiments; Renzi et al. (2014) also used a within-subjects design targeting with their stimulation the occipital face area, in this case only anodal and sham were used and experiments took place on two consecutive days. Looking at the same composite face task, here too it was found that in the sham condition the typical composite face effect is found with accuracy better for congruent than incongruent trials in the aligned condition while in the misaligned condition this is reduced. However, the key difference comes in the active condition where there is found no effect of tDCS on the composite face effect. Given the disparity in these results, stimulation delivered at PO8 offers an important contribution to the series of active control studies outlined in this section and extended by the tDCS experiments presented in chapters 2 and 6. Experiment 1 used an Fp3-Cz active control to investigate whether the tDCS induced effects on the inversion effect can be attributed solely to the placement of the anodal electrode at Fp3. In this case too, the results were found to be not significantly different from sham. Our results therefore support the findings from these previous studies indicating that it is the specific Fp3-Fp2 montage that modulates perceptual learning. Our active control Fp3-Cz set-up as was also used in Ambrus et al. (2011) but in this case was shown to significantly impair categorisation accuracy and eliminate the prototype effect, it may seem surprising then that in experiment 1 we find no significant difference in the Fp3-Cz montage and sham stimulation when investigating the inversion effect, but there are a number of potential reasons for these differing results. Firstly, the stimuli and tasks used differed, Ambrus et al. (2011) presented dot patterns in the context of a prototype distortion task whereas we used faces and checkerboards in a matching task, it could be that there are different processing mechanisms required to perform these tasks and as such they are differentially affected by tDCS stimulation. In addition to this, Ambrus et al. (2011) delivered their anodal stimulation at 1.0mA while in experiment 1 we delivered 1.5mA stimulation. The effects induced by tDCS do not necessarily increase linearly when stimulation intensity is increased, in fact in some cases increasing the intensity of stimulation can result in a reduction of such effects (Batsikadze et al., 2013; Esmaeilpour et al., 2018; Ehrhardt et al., 2021). When considering these crucial differences, it becomes clear why there is some incongruency in the effectiveness of tDCS in the Fp3-Cz set-up for these two experiments.

Although it was not specifically my intention to do so, the design of experiment 5 adds a final layer of specificity to the tDCS montage by showing that the effect of perceptual learning is not only determined by the positioning of one electrode on each of Fp3 and Fp2 but also that precise configuration of the anodal and cathodal electrodes on these areas is important. Reversing the positions of the electrodes also reverses the effect that stimulation has on perceptual learning. We can therefore support the notion that the reduction of the inversion effect is specific to the anodal Fp3-Fp2 tDCS montage. What it is specifically about this montage that is able to modulate perceptual learning is as yet not entirely clear but suggestions for future research involving fMRI that would enable this to be investigated are outlined in the future research section of this discussion.

7.3.2.1 Implications for the reversal of the perceptual learning effect with cathodal stimulation

While the existing body of literature has clearly demonstrated that anodal Fp3-Fp2 stimulation modulates perceptual learning and reduces the inversion effect through impairment for upright stimuli, we are not entirely sure by what process this modulation occurs. Civile, Verbruggen, et al. (2016) offer increased generalisation between stimuli as an explanation making it harder to discriminate similar exemplar from one another. This is supported by the predictions of the MKM model that perceptual learning increases discriminability by reducing generalisation of the common features and increasing the salience of the distinct feature. There were, however, alternative explanations that could also explain this effect. Changes in the way stimuli are encoded or the context in which this happens (e.g. via the use of different scan paths) may also be responsible either directly or indirectly for the tDCS-induced effect of perceptual learning. That we found in experiment 5 that cathodal Fp3-Fp2 stimulation is able to reverse the effect that anodal Fp3-Fp2 has on perceptual learning is good evidence to refute these latter explanations, however. Physical changes in how stimuli are encoded would have already occurred in the study phase and would likely not be reversed by mismatched encoding in the recognition phase. While there are still questions to be answered about how tDCS stimulation produces the effect that it does, this does at least aid our understanding by eliminating these particular explanations. Further evidence about the specific impact that cathodal tDCS is having in this paradigm could be provided by replication of experiment 5 with the use of checkerboard stimuli in addition to the face stimuli, based on the previous tDCS work and the interpretation in terms of the MKM model we would expect for this effect to carry over on to the checkerboard stimuli, reversing the decrease in perceptual learning and increase in generalisation that is thought to impair performance for familiar upright stimuli.

7.3.2.2 Neural correlates of the inversion effect and extension to cathodal tDCS stimulation

Work aiming to characterise the effect of anodal tDCS on perceptual learning has been conducted in a number of EEG studies. There is an extensive body of work showing that the face inversion effect can be seen not only behaviourally in terms of recognition accuracy but also in the ERP components associated with face stimuli. The N170 ERP component peaks around 130-220ms after the onset of an upright face stimulus, while for inverted faces the N170 is delayed but often greater in terms of amplitude (Rossion et al., 2000). While there is not yet a definitive explanation of this effect, it is often interpreted to be the results of a disruption to configural processing due to inversion and this is supported by examination of the N170 for faces manipulated by scrambling or Thatcherisation which shows a reduced inversion effect (reduced delay in the N170 for upright and inverted stimuli). In a similar way to behavioural inversion effects, the differential pattern of results on the N170 has been linked to familiarity with prototype-defined stimuli; participants trained to recognise categories of Greebles showed a smaller N170 inversion effect for Greebles than faces pre-training but post-training the inversion effect was comparable between the two with inverted stimuli showing a delayed peak compared to upright (Rossion et al., 2002). Similarly, Civile, Zhao et al.'s (2014) work with checkerboards (explored in depth in chapter 1) included an EEG study which showed that for prototype categories of checkerboards which had been trained in the categorisation phase (familiar) there was a larger inversion effect (delayed N170 with a greater amplitude) in the subsequent recognition task than was present for prototype categories not seen previously (novel).

In addition to these experiments focusing on behavioural paradigms the N170 has also been examined in the context of tDCS experiments. Civile, Waguri, et al. (2020) used an anodal Fp3-Fp2 tDCS procedure in combination with EEG during an old/new recognition task with faces and found that, as with behavioural manipulations, anodal tDCS resulted in a significantly reduced inversion effect compared to sham, shown in this context by a lesser difference in delay in the N170 for the upright relative to inverted faces. This falls in line with the explanation that modulation of the N170 is reflective of the use of perceptual expertise to utilise configural information; in the sham condition there is a large inversion effect as participants are able to use this perceptual expertise for the upright faces but not the inverted while in the anodal condition perceptual learning is eliminated and thus this advantage for upright faces is lost, resulting in a reduction in the inversion effect. Further exploration on the effect of cathodal stimulation on the N170 ERP component would allow us to further characterise the effect that it has and contribute to the body of literature seeking to explain how tDCS is able to modulate perceptual learning. Given that in experiment 5 we see the typical reduction in the inversion effect (in terms of accuracy) under anodal stimulation and the reversal of this effect back to sham level when cathodal stimulation is applied following anodal stimulation, we would expect, following the perceptual learning explanation, that the N170 would follow a similar pattern. During anodal Fp3-Fp2 stimulation (as previously seen in Civile, Waguri, et al., 2020) there is a reduction in the N170 inversion effect in terms of latency as perceptual learning for upright faces is impaired, as this behavioural perceptual learning effect is reversed under cathodal stimulation we would expect to see that the inversion effect for the N170 also increases again, with a greater delay again seen for the inverted faces.

7.4 Contour manipulation experiments

7.4.1 Implications for holistic processing (Hole, et al., 1999; Civile, et al. 2016)

The theory underlying the contour manipulations used in experiments 2, 3a, 3b, and 4 is based on that of holistic processing outlined in Hole, et al. (1999) and consolidated in Civile, et al. (2016). This states relational processing can be divided into two distinct types: configural information which refers to the specific spatial relationship between the features of the face and the holistic processing which is much broader and will be evoked for anything that conforms to the basic shape of a face. That there is a distinction between types of relational processing is supported by Hole, et al.'s work with photographic negatives, these have been shown to impair recognition performance for faces similarly to inversion but without the disruption to configuration processing thought to play a role in the inversion effect. That negatives are indeed subject to configural processing is demonstrated by the finding that they produce the chimeric face effect in much the same way as positives do. Negatives must therefore be targeting a different aspect of face recognition when reducing overall performance. Hole, et al. (1999) cite holistic processing as a possible explanation and offer face contour and first-order information as the potential source of this. Civile, et al. (2016) supported this with the discovery that when single-feature information (proposed to be salient in face recognition) is controlled for and second-order information is disrupted in their New Thatcherised faces, first-order information still produces a significant inversion effect. The findings from experiment 2 offer further support for face contour as a component of holistic processing, as contour manipulation in scrambled faces is able to reduce the usually robust inversion effect found for these stimuli, which indicates the perhaps there is some additional disruption to holistic information driving the effect. This notion is to an extent also supported by the findings in experiment 3a and 3b, particularly when low performers are excluded or the data is pooled.

7.4.2 Face contour and overall performance

Experiment 4 offers evidence that while the inversion effect was not significantly reduced for New Thatcherised faces with the contour manipulation, there was a significant reduction in overall recognition performance. This was also the case in the original analysis of normal faces. Taken together these findings indicate additional support for the role of face contour (whether as a component of holistic processing or otherwise) in face recognition. This conclusion is hampered somewhat by the lack of an overall effect on performance for the scrambled faces. As discussed in chapter 5, floor level recognition performance has been refuted as an explanation for this by the additional analysis in which performance was significantly above chance for all stimulus conditions. However, this does not preclude difficulty as an explanation all together, while significantly above chance it may still be the case the recognition for manipulated scrambled faces is more difficult that either that for manipulated New Thatcherised faces or manipulated normal faces (particularly for the inverted stimuli which are already more difficult to recognise). Simple observation of the stimuli tells you that the scrambled faces are distorted to a greater extent than the other two stimulus types and therefore this theory seems plausible. Improving baseline performance using a matching task rather than an old/new recognition task was the focus of a pilot experiment I conducted on the scrambled faces. In the same way as in experiment 2 scrambled normal contour and scrambled blurred contour faces were compared in a within-subjects design. The matching task involved the presentation of a face (upright or inverted) followed by an interval mask and then a second face, with participants required to make a same/different judgement on the second face in comparison to the first. In this case the was no significant reduction in the inversion effect found for blurred contour faces compared to normal (in fact blurred contour inverted face were numerically higher than blurred contour upright faces) but there was a significant reduction in overall performance as a result of the blurring manipulation. This provides preliminary support for the findings in experiment 3a, 3b, and 4 in relation to the effect of face contour on face recognition performance.

There is an argument to be made here that the vastly increased performance made possible in a matching task may mean that any impact on the inversion is more difficult to detect and that the matching task is simply not sensitive enough to make any differential effect on upright and inverted faces observable. On the one hand, we have found analogous results in the tDCS literature with old/new recognition tasks and matching tasks, indicating that in this type of experiment at least they are comparable. However, it seems plausible that there may be distinct differences between the effects found in a tDCS experiment compared to a purely behavioural paradigm, meaning that a matching task might still struggle to show less prominent differences between upright and inverted stimuli for a behavioural manipulation. This finding nevertheless has interesting implications, if we assume that the matching task provides a reliable measure of the effect of contour manipulation then a more extensive replication of this pilot and extension to the other face stimuli used in experiments 3a, 3b, and 4 could provide key evidence about the role of holistic processing in face recognition.

7.4.3 The chosen contour manipulation

The use of a contour manipulation for the purposes of investigating holistic processing has clear advantages in that the stimuli it produces conform to the basic shape of a face to a much lesser extent (at least visually) than faces with unaltered contour. It also leaves the configural information between the main features of the face unaltered i.e., the spatial relationship between the eyes, nose, and mouth is not affected. However, there is also a clear limitation of this manipulation; the face outline itself has a spatial relationship with the internal features of the face and thus shares configural information with them which is disrupted by the manipulation. The spiked outlines help somewhat in the sense that a portion of the configural information that makes the faces distinct is used to create the outlines and therefore some of the differences in information are preserved relative to other exemplars with the spiked outline (although in a different form which does not offer the expertise benefit usually conferred by the configural information). However, there is still some configural information that is lost as a result of changing the shape of the outline and all of the configural information relating to the outline is different compared to the faces pre-manipulation. Future work that would address this and thus help to clarify the role of the face outline in face recognition, would be to create a set of faces for which the outline is not manipulated but the features inside are pixelated. This is not dissimilar to the work done by Davidenko (2007) which showed that silhouetted face profiles can produce a significant inversion effect. The production of a significant inversion effect for outlines with pixelated features would offer strong evidence in support of the face outline as holistic information.

7.5. Future Research

7.5.1 Combined tDCS and fMRI

The findings from experiment 1 contribute to a body of literature supporting the notion that there is something specific about the Fp3-Fp2 montage used across a wide range of tDCS studies which results in the modulation of perceptual learning while none of the variations of active control were able to achieve the same result. What specifically it is about this montage that causes this effect is not yet clear and the use of fMRI along with tDCS experiments may help us to develop a better understanding of this. There have already been some studies that use fMRI in conjunction with tDCS but without the specific context of inversion effect tasks. Park et al. (2013) investigated the effect of tDCS over the left DLPFC using resting state fMRI, Anodal and sham condition were used (between-subjects) with active stimulation delivered at 1.0mA for 20 minutes. fMRI data were collected pre and post stimulation for both the anodal and sham conditions and participants were instructed to remain motionless with their eyes closed throughout. The results showed that interhemispheric connectivity to the left DLPFC increased following tDCS compared to sham and reached the frontal, temporal and subcortical regions of the right hemisphere. Connectivity to the frontal regions around the left DLPFC however decreased following tDCS stimulation. Peña-Gómez et al. (2012) similarly acquired resting state fMRI data following tDCS to the left DLPFC, in this case a within-subjects design was used and participants underwent two days of experimentation, one involving sham stimulation followed by anodal stimulation and the other involving sham followed by cathodal stimulation. Their aim was to investigate the functioning of different brain networks, the default-mode network (associated with resting states and control conditions) and the anti-correlated network (associated with cognitive processing, particularly when focused attention to external cues is required). The DLPFC forms part of the AN and this was the reason for targeting stimulation to this area. Findings showed that either form of active stimulation increased connectivity between prefrontal and parietal regions (also part of the AN) while the spatial robustness of the DMN was reduced compared to sham. Interestingly, both these studies both explore their results in the context of explaining why for some tasks anodal tDCS is able to improve performance. While these studies offer some insight into how tDCS modulates connectivity between brain regions it is important to note that they do so in resting states, combining this body of research with the face recognitions tasks may offer some interesting additional insights into how this stimulation affects the functioning of brain areas as relates to specific cognitive tasks. There are few different variations of this study that may be of relevance in the future; exploration of the changes that may occur for the Fp3-Fp2 montage compared to sham and active control may provide a clearer understanding of how this specific montage is able to modulate perceptual learning while others do not. Additionally, comparison of this effect during face and checkerboard recognition tasks could allow us to assess whether there are functional differences in how tDCS acts on these different stimuli, this is particularly of interest given the slightly differential effects of tDCS on the inversion effect (entirely eliminated for checkerboards and reduced but still significant for faces). Using fMRI to operationalise the different effects of anodal and cathodal stimulation during face recognition tasks may also offer some notable results, particularly given that Peña-Gómez et al. (2012) found that both forms of active stimulation resulted in similar changes to functional connectivity.

7.5.2 The effect of cathodal stimulation alone on perceptual learning

We have shown in experiment 5 that cathodal Fp3-Fp2 stimulation can reverse the negative impact of anodal Fp3-Fp2 stimulation on perceptual learning, at least back up the level it was before tDCS (i.e., in line with sham). An interesting question for future research to answer is whether delivering cathodal Fp3-Fp2 stimulation without first disrupting perceptual learning would increase the perceptual learning effect and lead to a greater inversion effect than sham by improving performance for upright faces. Generally, humans already have an incredible ability to recognise faces and do so in everyday life with great accuracy so it is unclear whether it would be possible improve his any further but none-the-less it is a possibility worth exploring in future research. In addition, the use of cathodal stimulation during checkerboard recognition tasks would allow a similar kind of comparison to that seen in experiment 1, if the effects of anodal tDCS can be reversed for checkerboards in a similar way as has been observed faces this would help support the concept of a shared processing mechanism for both stimulus types (namely perceptual learning) which can be modulated by tDCS.

A further interesting application of the cathodal Fp3-Fp2 stimulation would be to test it on participants with prosopagnosia. Often referred to as face blindness, prosopagnosia refers to a disorder resulting in impaired recognition of individual faces, the severity can range from mild difficulty in remembering faces to complete inability to discriminate between any faces, and it is usually acquired as a result of brain damage, and unrelated to intellectual deficiency or visual impairment (e.g., Bodamer, 1947; Rondot & Tzavaras, 1969). There is already evidence that training perceptual learning in prosopagnosics can improve face recognition, with improvement maintained at 3-month follow up (Corrow et al. 2019), and so it certainly seems possible that cathodal Fp3-Fp2 tDCS could aid this process, at least in facilitating better training if not as a stand-alone therapy due to the relatively short-term effects of tDCS.

7.5.3 Level of perceptual expertise for checkerboards

The results from chapter 2 experiment 1 have been explained in terms of the remaining inversion effect for faces compared to checkerboards in the anodal condition being due to a face-specific process that is unaffected by the tDCS.

However, we must also consider that the level of expertise gained for checkerboards during the categorisation task does not come close to the level of expertise gathered from our lifelong experience with faces. It is possible this lower degree of perceptual learning for checkerboards is simply easier to eradicate with anodal stimulation than the immense amount of perceptual learning that has taken place for faces. One potential avenue for future research would be to increase the level of expertise that participants are able to gain for checkerboards. Although no artificial intervention is ever going to match this to the extent that we have expertise for faces, it would at least allow us to assess whether greater opportunity for perceptual learning about checkerboards leads to a less reduced inversion effect than is currently obtained during tDCS. It would of course be hugely impractical for participants to gain very much expertise in a laboratory setting but the use of at home programmes could be used in place of this. An app would allow participants to experience checkerboards in a range of different categorisation and recognition paradigms while also allowing the researchers to track factors such as the amount time spent on this and participants' accuracy for these tasks. Following the achievement of a set performance threshold researchers could then apply the tDCS procedure and assess whether i) a significant inversion effect can be maintained for checkerboards during anodal stimulation as is the case for faces or ii) the difference in the inversion effect between anodal and sham differs for participants who have gained a greater level of perceptual expertise with checkerboards. Significant findings on either of the measure would suggest that perhaps the disparity in the level of expertise for faces and checkerboards contributes to the results from experiment 1 and Civile, Quaglia, et al. (2021).

7.6 Overall Summary

In conclusion, the research presented in this thesis has addressed one of the long-standing debates in the face recognition literature which has primarily focused on two opposing accounts: the first began by Yin (1969) positing that the face inversion effect (the reduction in recognition performance for inverted faces compared to upright) is the result of a face specific mechanism and the second from Diamond and Carey (1986) arguing that the face inversion effect results from expertise with prototype-defined stimuli. This thesis followed a line of research into perceptual learning as a specific process of expertise and extended comparisons of the inversion effect for faces and checkerboards when perceptual learning is disrupted by tDCS. The main findings in this regard are 1) expertise and facial specificity both play a role in the face inversion effect; 2) holistic processing may be one aspect of this face specific component and can be elicited by the face contour and first-order configural information; 3) the tDCS induced effects on the inversion effect are produced specifically by an anodal Fp3-Fp2 montage with no other active control condition obtaining similar results; 4) reversing the polarity of stimulation can reverse the reduction in the inversion effect resulting from anodal tDCS, indicating that disruption to perceptual learning better explains these tDCS induced effects than alternative theories.

REFERENCES

Ambrus, G. G., Zimmer, M., Kincses, Z. T., Harza, I., Kovács, G., Paulus, W., & Antal, A. (2011). The enhancement of cortical excitability over the DLPFC before and during training impairs categorization in the prototype distortion task. *Neuropsychologia*, *49*(7), 1974-1980.

Batsikadze, G., Moliadze, V., Paulus, W., Kuo, M. F., & Nitsche, M. (2013). Partially non-linear stimulation intensity-dependent effects of direct current stimulation on motor cortex excitability in humans. *The Journal of physiology*, *591*(7), 1987-2000.

Busigny, T., & Rossion, B. (2010). Acquired prosopagnosia abolishes the face inversion effect. *Cortex*, *46*(8), 965-981.

Carey, S., & Diamond, R. (1994). Are faces perceived as configurations more by adults than by children?. *Visual cognition*, *1*(2-3), 253-274.

Civile, C., Cooke, A., Liu, X., McLaren, R., Elchlepp, H., Lavric, A., & McLaren, I. P. L. (2020). The effect of tDCS on recognition depends on stimulus generalization: Neuro-stimulation can predictably enhance or reduce the face inversion effect. *Journal of Experimental Psychology: Animal Learning and Cognition*, *46*(1), 83

Civile, C., McLaren, R. P., & McLaren, I. P. (2014). The face inversion effect—Parts and wholes: Individual features and their configuration. *Quarterly journal of experimental psychology*, *67*(4), 728-746.

Civile, C., McLaren, R., & McLaren, I. P. (2016). The face inversion effect: roles of first and second-order configural information. *The American Journal of Psychology*, *129*(1), 23-35.

Civile, C., McLaren, R., & McLaren, I. P. L. (2018). How we can change your mind: Anodal tDCS to Fp3 alters human stimulus representation and learning. *Neuropsychologia*, *119*, 241-246

Civile, C., McLaren, R., Waguri, E., & McLaren, I. P. (2020). Testing the immediate effects of transcranial Direct Current Stimulation (tDCS) on face recognition skills. In *CogSci*.

Civile, C., McLaren, R., Milton, F., & McLaren, I. P. L. (2021). The effects of transcranial direct current stimulation on perceptual learning for upright faces and its role in the composite face effect. *Journal of Experimental Psychology: Animal Learning and Cognition*, *47*(1), 74.

Civile, C., McLaren, R., Cooke, A., & McLaren, I. P. L. (2022). Using transcranial direct current stimulation (tDCS) to influence decision criterion in a target detection paradigm. *Journal of Experimental Psychology: Animal Learning and Cognition*.

Civile, C., McLaren, R., Forrest, C., Cooke, A., & McLaren, I. P. (2023). Modulating perceptual learning indexed by the face inversion effect: Simulating the application of transcranial direct current stimulation using the MKM model. *Journal of Experimental Psychology: Animal Learning and Cognition*, *49*(3), 139.

Civile, C., Quaglia, S., Waguri, E., Ward, M., McLaren, R., & McLaren, I. P. L. (2021). Using transcranial Direct Current Stimulation (tDCS) to investigate why faces are and are not special. *Scientific reports*, *11*(1), 4380.

Civile, C., Verbruggen, F., McLaren, R., Zhao, D., Ku, Y., & McLaren, I. P. L. (2016). Switching off perceptual learning: Anodal transcranial direct current stimulation (tDCS) at Fp3 eliminates perceptual learning in humans. *Journal of Experimental Psychology: Animal Learning and Cognition*, *42*(3), 290.

Civile, C., Wooster, B., Curtis, A., McLaren, R., McLaren, I. P., & Lavric, A. (2019). Using transcranial Direct Current Stimulation (tDCS) to modulate the face inversion effect on the N170 ERP component. Cognitive Science Society.

Civile, C., Zhao, D., Ku, Y., Elchlepp, H., Lavric, A., & McLaren, I. P. (2014). Perceptual learning and inversion effects: Recognition of prototype-defined familiar checkerboards. *Journal of Experimental Psychology: Animal Learning and Cognition*, *40*(2), 144.

Civile, C., Waguri, E., Quaglia, S., Wooster, B., Curtis, A., McLaren, R., ... & McLaren, I. P. L. (2020). Testing the effects of transcranial direct current stimulation (tDCS) on the face inversion effect and the N170 event-related potentials (ERPs) component. *Neuropsychologia*, *143*, 107470.

Cunillera, T., Fuentemilla, L., Brignani, D., Cucurell, D., & Miniussi, C. (2014). A simultaneous modulation of reactive and proactive inhibition processes by anodal tDCS on the right inferior frontal cortex. *PloS one*, *9*(11), e113537.

Davidenko, N. (2007). Silhouetted face profiles: A new methodology for face perception research. *Journal of Vision*, *7*(4), 6-6.

de Gelder, B., & Rouw, R. (2000). Paradoxical configuration effects for faces and objects in prosopagnosia. *Neuropsychologia*, *38*(9), 1271-1279.

Diamond, R., & Carey, S. (1986). Why faces are and are not special: an effect of expertise. *Journal of experimental psychology: general*, *115*(2), 107.

Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in psychology*, *5*, 781

Ellis, H. D., Shepherd, J. W., & Davies, G. M. (1979). Identification of familiar and unfamiliar faces from internal and external features: Some implications for theories of face recognition. *Perception*, *8*(4), 431-439.

Ehrhardt, S. E., Filmer, H. L., Wards, Y., Mattingley, J. B., & Dux, P. E. (2021). The influence of tDCS intensity on decision-making training and transfer outcomes. *Journal of Neurophysiology*, *125*(2), 385-397.

Esmaeilpour, Z., Marangolo, P., Hampstead, B. M., Bestmann, S., Galletta, E., Knotkova, H., & Bikson, M. (2018). Incomplete evidence that increasing current intensity of tDCS boosts outcomes. *Brain stimulation*, *11*(2), 310-321.

Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision research*, *35*(14), 2089-2093.

Gauthier, I., & Tarr, M. J. (1997). Becoming a “Greeble” expert: Exploring mechanisms for face recognition. *Vision research*, *37*(12), 1673-1682.

Ghent, L., Bernstein, L., & Goldweber, A. M. (1960). Preferences for orientation of form under varying conditions. *Perceptual and Motor Skills*, *11*(1), 46-46.

Haig, N. D. (1984). The effect of feature displacement on face recognition. *Perception*, *13*(5), 505-512.

Henle, M. (1942). An experimental investigation of past experience as a determinant of visual form perception. *Journal of Experimental Psychology*, *30*(1), 1.

Hole, G. J., George, P. A., & Dunsmore, V. (1999). Evidence for holistic processing of faces viewed as photographic negatives. *Perception*, *28*(3), 341-359.

Hosie, J. A., Ellis, H. D., & Haig, N. D. (1988). The effect of feature displacement on the perception of well-known faces. *Perception*, *17*(4), 461-474.

Jacobson, L., Javitt, D. C., & Lavidor, M. (2011). Activation of inhibition: diminishing impulsive behavior by direct current stimulation over the inferior frontal gyrus. *Journal of cognitive neuroscience*, *23*(11), 3380-3387.

Köhler, W. (1960). *Dynamics in psychology*. WW Norton & Company.

Leder, H., & Bruce, V. (1998). Feature processing from upright and inverted faces. In *Face recognition: From theory to applications* (pp. 547-555). Berlin, Heidelberg: Springer Berlin Heidelberg.

Levine, D. N., & Calvanio, R. (1989). Prosopagnosia: A defect in visual configural processing. *Brain and cognition*, *10*(2), 149-170.

Lewis, M. B., & Johnston, R. A. (1997). The Thatcher illusion as a test of configural disruption. *Perception*, *26*(2), 225-227.

McKone, E., & Yovel, G. (2009). Why does picture-plane inversion sometimes dissociate perception of features and spacing in faces, and sometimes not? Toward a new theory of holistic processing. *Psychonomic bulletin & review*, *16*(5), 778-797.

McLaren, I. P. L. (1997). Categorization and perceptual learning: An analogue of the face inversion effect. *The Quarterly Journal of Experimental Psychology Section A*, *50*(2), 257-273.

McLaren, I. P. L., Carpenter, K., Civile, C., McLaren, R., Zhao, D., Ku, Y., & Verbruggen, F. (2016). Categorisation and perceptual learning: Why tDCS to Left DLPC enhances generalisation. *Associative learning and cognition: Homage to Professor NJ Mackintosh, 37-67.*

McLaren, I. P., Kaye, H., & Mackintosh, N. J. (1989). An associative theory of the representation of stimuli: Applications to perceptual learning and latent inhibition.

Nitsche, M. A., & Paulus, W. (2000). Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. *The Journal of physiology*, *527*(Pt 3), 633.

Park, C. H., Chang, W. H., Park, J. Y., Shin, Y. I., Kim, S. T., & Kim, Y. H. (2013). Transcranial direct current stimulation increases resting state interhemispheric connectivity. *Neuroscience letters*, *539*, 7-10.

Peña-Gómez, C., Sala-Lonch, R., Junqué, C., Clemente, I. C., Vidal, D., Bargalló, N., & Bartrés-Faz, D. (2012). Modulation of large-scale brain networks by transcranial direct current stimulation evidenced by resting-state functional MRI. *Brain stimulation*, *5*(3), 252-263.

Peterson, M. F., & Eckstein, M. P. (2013). Individual differences in eye movements during face identification reflect observer-specific optimal points of fixation. *Psychological science*, *24*(7), 1216-1225.

Rakover, S. S., & Teucher, B. (1997). Facial inversion effects: Parts and whole relationship. *Perception & Psychophysics*, *59*(5), 752-761.

Renzi, C., Ferrari, C., Schiavi, S., Pisoni, A., Papagno, C., Vecchi, T., ... & Cattaneo, Z. (2015). The role of the occipital face area in holistic processing involved in face detection and discrimination: A tDCS study. *Neuropsychology*, *29*(3), 409.

Scapinello, K. F., & Yarmey, A. D. (1970). The role of familiarity and orientation in immediate and delayed recognition of pictorial stimuli. *Psychonomic Science*, *21*(6), 329-330.

Seger, C. A., Poldrack, R. A., Prabhakaran, V., Zhao, M., Glover, G. H., & Gabrieli, J. D. (2000). Hemispheric asymmetries and individual differences in visual concept learning as measured by functional MRI. *Neuropsychologia*, *38*(9), 1316-1324.

Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British journal of psychology*, *75*(2), 221-242.

Stramaccia, D. F., Penolazzi, B., Sartori, G., Braga, M., Mondini, S., & Galfano, G. (2015). Assessing the effects of tDCS over a delayed response inhibition task by targeting the right inferior frontal gyrus and right dorsolateral prefrontal cortex. *Experimental brain research*, *233*, 2283-2290.

Tanaka, J. W., & Farah, M. J. (1991). Second-order relational properties and the inversion effect: Testing a theory of face perception. *Perception & Psychophysics*, *50*(4), 367-372.

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, *46*(2), 225-245.

Tanaka, J. W., & Sengco, J. A. (1997). Features and their configuration in face recognition. *Memory & cognition*, *25*(5), 583-592.

Thompson, P. (1980). Margaret Thatcher: A new illusion. *Perception*, *9*(4), 483-484.

Valentine, T., & Bruce, V. (1986). The effect of race, inversion and encoding activity upon face recognition. *Acta psychologica*, *61*(3), 259-273.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of experimental psychology*, *81*(1), 141.

Yovel, G., & Kanwisher, N. (2005). The FFA shows a face inversion effect that is correlated with the behavioral face inversion effect. *Journal of Vision*, *5*(8), 632-632.