THE UNIVERSITY
of EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

Multisensory Integration, Predictive Coding and the Bayesian Brain: Reintegrating the Body Image and Body Schema Distinction into Cognitive Science.

Ashleigh Watson

Doctor of Philosophy   The University of Edinburgh   2016

Declaration

This is to certify that that the work contained within has been composed by me and is entirely my own work. No part of this thesis has been submitted for any other degree or professional qualification. Signed:

Abstract

The classic distinction between the body schema and the body image received renewed interest in cognitive psychology, in part because of the attempts by the leading psychologist Charles Spence and his co-authors to synthesise a mounting body of research into the multisensory nature and functional properties of the neural structures in primate cortex that are sensitive and responsive to cross-modal stimuli generated from the body and objects located close to the body, and the famous rubber hand illusion which purported to illustrate how the perception and understanding of what counts as one's body, i.e., our body image, can be manipulated to include foreign, body-part-like, objects such as a rubber hand. This approach was intended to settle age old questions about how the body schema – the system sub-personal sensorimotor system that shapes, facilitates and regulates motor control – is implemented in the brain and address historic confusions about how the body schema should be understood as an explanatory concept, as well as the problems surrounding the body schema and image distinction on the grounds of the persistent conflation between the two concepts. However, after offering several proposals as to how the body schema should be used to organise and interpret the empirical data, the distinction fell out of favour with Spence and his colleagues on the grounds of the very problems they intended to resolve. The proposed solution is an alternative theoretical framework that, I shall argue, never materialised. Instead, the various definitions they disseminate, I will claim, simply serve to further perpetuate the same problems and confusions about the body schema. Thus, the current state of the literature on the body image and schema in cognitive psychology is in dire need of a conceptual framework that would help us situate and interpret the important empirical data. I propose that we revisit the philosophical debates that were inspired by the philosopher Shaun Gallagher as part of his project to provide a conceptual analysis of the body schema and image distinction and vindicate its status as an important explanatory device for the explanatory ambitions of embodied cognition. Gallagher's analysis opens up important questions about how the sub-personal multisensory processes of the body schema not only facilitate moment-by-moment motor behaviours, but how they shape and *optimise* motor control across developmental timelines, as well the importance of the embodied configuration of an agent and its particular eco-niche for shaping and facilitating its motor behaviours. The second important argument of the thesis is that the response to Gallagher's analysis has simply served to suppress the line of research that Gallagher inspired because the questions his analysis raises have been overshadowed by more general disputes between

Gallagher and his opponents about the shape an analysis of the body schema from the perspective of embodied cognition should take. As such, potentially promising lines of research in relation to the body schema have since dried up. As part of my attempt to make progress on the issues that are laid out at the first and second stages of the thesis, the third stage will involve an exploration into the seminal Bayesian approach to understanding cross-modal cue optimisation as it applies to object perception (Banks & Ernst, 2002) and the recent extension of this paradigm to the multimodal sensorimotor processes that underpin motor behaviour in action-oriented cognitive science (e.g., Friston, 2010). The conclusion of the thesis is that the move from an embodied to an action-oriented analysis of the body schema, and the conceptual distinction of which it is part, provides us with the right kind of theoretical resources to begin to pursue fruitful avenues of research that allow us to begin to address the questions set out by Gallagher's analysis whilst avoiding (some of) the pitfalls that beset the embodied approach. In the final chapter I use this model of the body schema to illustrate how it can provide the basis for working back up towards a comprehensive theory of the body image and schema distinction, which I then bring to bear on current, as-yet-unaddressed, issues in developmental psychology.

Lay Summary

To effectively move our body around in our surroundings we (or rather our brain) must continually keep track the changing spatial relations between the body and objects in space to prevent harmful collisions and work out how we need to move our body in to interact with objects of interest. To achieve this endeavour, our brain is tasked with the ceaseless job of processing and utilising the various sources of sensory information that it is constantly bombarded with, e.g., from vision, audition, touch and proprioception (i.e., information from our body about its postural configuration and the position of our limbs). A key question that is of interest to psychologists is how the brain uses this information for the purposes just mentioned. Does it just use one source of information, several, or all? What happens if two or more sensory channels provide conflicting information? Such questions shape the research area known as 'multisensory integration' which involves the collaborations between philosophers, psychologists, roboticists, and the like to understand the significance of multisensory collaboration for our mental and cognitive lives. Perhaps the most interesting challenge for multisensory integration theorists is to understand and explain why, and how, some of the multisensory information that is processed by the brain enters our conscious awareness, and some does not. Specifically, we will look at the questions why and how the brain synthesises multisensory information to produce motor responses without us being aware of it. For instance, as you are driving along in your car thinking about what you're going to have for tea, your body will be making all sorts of adjustments. For instance, your hands might subtly guide the steering wheel as you travel along the road in very subtle ways. Likewise, as you reach your hand downwards to shift gears you may be unaware of the fact that, at some point, your fingers have formed into a grasp to clutch the gear stick. To distinguish the aspects of your experience of the movement of your body of which you are aware from those you are not, neurologists, philosophers and psychologists make a distinction between the 'body image' and the 'body schema', respectively. The body image categorises the multisensory processes pertaining to body perception and the body schema pertains to the multisensory processes that facilitate our motor responses outwith the scope of attention and awareness. However, this distinction has a long and troubled history because of the regular conflation between the two concepts. As such, when theorists talk of the body schema and body image they are often talking at crossed purposes and this has long stood in the way of collaborative research. The purpose of the first stage of this thesis is to show that

these problems still thrive in the literature. The second stage of the thesis is to make use of current developments in cognitive science pertaining to how multisensory processing facilitates perceptual and motor processes. The third stage is to show how we can use this model to resolve the problems surrounding the body image/body schema distinction and how this new-and-improved analysis of the distinction can help developmental psychologists who are interested seek an analysis of this distinction to assist their developmental studies into the developmental significance of multisensory processing for body experience and movement.

Acknowledgments

I'd like to take this opportunity to thank, first and foremost, my primary supervisor Dave Ward for his helpful advice, comments for improvement, and most of all for his patience. I also owe him a great deal of gratitude for directing me towards the work of Shaun Gallagher at the initial stages of writing, after spending many months running around in circles trying to figure out what on earth the term 'body schema' was supposed to mean. I'd also like to thank my secondary supervisor Suilin Lavelle for her helpful commentary and advice. I also owe thanks to some of my Edinburgh PhDs for their instructive, informal and friendly feedback at various reading groups and work-in-progress seminars. I'd like to thank my parents, Michael and Jacqueline, for their indefatigable support, love and endless supply of coffee. I'd also like to thank my grandmother, Grace Furey, and my grandfather, John "Jack" Furey for their encouragement and support. Finally, I dedicate this in memory of my wonderful grandmother, Marion, and my loyal chum Molly, whose canine companionship I missed sorely at the final stages of writing.

Introduction

This key objective of this thesis is to revisit and re-evaluate concurrent discussions that took place in cognitive psychology and philosophy about the optimal significance of multisensory integration for motor processes around a decade ago. Broadly speaking, the guiding question is why information from multiple sources of sensory information from the body and environment is required for motor output. In psychology, the question is how the brain represents the body and its relation to objects within reachable distance (also called 'peripersonal space'). The rationale is that objects which are closer to the body require different motoric responses relative to those which are farther away. For example, as I'm walking along the pavement and overhanging tree branch will not cause a change in my motor responses until my I get relatively close to it, at which point I may duck my head. Likewise, objects that are closer to the body offer different opportunities for interaction compared to objects that are farther away. For example, I can grasp the cup that's in front of me by reaching my arm out towards it. However, if it is in a separate room I must relocate my entire body by walking to, say, the kitchen. As we will see in chapter one, the discovery of neural structures in non-human primate cortex showed that there are a population of neurons in regions of the brain known to be involved in localising the body in space and motor planning, which are simultaneously responsive to tactile and proprioceptive (i.e., the sensory modality responsible for providing information about the spatial position and movement of the body parts) information from the body and to visual information from objects that are within very close range of the body. Such discoveries, the story goes, provide empirical leverage to the existence of the so-called 'body schema'. The body schema is an explanatory concept which refers to the system of processes involved in localising the body parts in space to regulate the motor responses of the body. Importantly, the body schema is a sub-personal system that operates below the threshold of consciousness. The body schema is one half of an important conceptual distinction (along with the body image) that had, by the time this important body of work was taking place, repeatedly fallen in and out of favour with cognitive psychologists on the grounds of the persistent conflation between the two concepts. Thus, the discovery of this class of neurons holds promise of localising the neural underpinnings of the so called 'body schema' and tells us something additional about their nature: that its sensory processes are *multisensory* and are also responsive to objects close to the body. The significance of this, the authors claim, is that it goes some way towards

resolving the age-old confusions which surrounded the distinction. However, despite the seemingly straightforward implications of the empirical data for our understanding of the body schema, the direction in which the authors take the concept only serves to further perpetuate the conceptual confusions already associated with the concept, and puts the future of the distinction in jeopardy, or so I will argue in chapter one.

At around the same time the philosopher Shaun Gallagher established his final attempt at providing a clear conceptual dissociation between the concept of the body image and the body schema and provided an analysis of the functional components that make up each system and provided an analysis of how the two systems work in conjunction in the context of everyday action (amongst many other things). In a similar fashion to the approach in cognitive psychology just outlined, Gallagher defines the body schema in terms of a sub-personal system of multisensory sensorimotor processes that are involved in producing and regulating motor behaviour. In contrast, the body image is comprised out perceptual, conceptual and affective processes that shape our personal-level awareness and apprehension of our body. Therefore, the distinction is, for the most part, functional because Gallagher outlines the components in terms of what they do. Within his analysis, Gallagher throws up important questions about how the body schema and body percept, individually and collectively, optimise motor performance. To my mind, the importance of his analysis has gone unrecognised for the one response to his framework from Vignemont mounts an objection against Gallagher's conception of the body schema because he doesn't conceptualise the body schema in terms of mental representation in the brain. I will argue that this objection is unwarranted, and that her alternative model perpetuates the very confusions that Gallagher was trying to eradicate. However, I will claim that there is an important line of thought that Vignemont raises but does not explore: the importance of prediction for action.

This provides the basis for chapters 3 and 4 of the book which explore two mathematically inspired models of the optimising strategies of sensory evaluation and integration. In chapter 3 we will explore a seminal study from Marc Ernst which shows that visuo-tactile inputs are brought together to provide a calibrated estimate of a perceptual input. This is based on the idea, taken from Bayesian models of neural functioning, that the brain is not just a reactive mechanism which recapitulates the inputs it receives as outputs. Rather the brain actively integrates cross-modal inputs that carry equivalent information to maximise the accuracy (i.e., optimise) of its outputs. This is based on the principle that the reliability of the information provided by each sensory modality is task- and context- dependent. To evaluate the inputs in the specified the brain brings any relevant prior-knowledge and beliefs (i.e.,

'priors') to bear on the incoming sensory data to provide the best possible output. Importantly, this Bayesian model takes perceptual processes to be *probabilistic* in nature. Thus, the outcome of perceptual processes is the best possible guess about the source of the inputs. Then in chapter 4 we will consider a recent Bayesian inspired model which extends Ernst's story about perceptual processes to sensorimotor processes. This is the prediction error minimisation theory. This model accepts the Bayesian tenets just outlined, however one the distinctive hallmarks of this model is the claim that the bulk of the knowledge-driven estimation processing occurs *prior* to the stimulus entering the low-level sensory organs. Thus, instead of estimating the cause of the inputs based on the information that is received, the optimising strategies of the brain estimates the likely incoming sensory data before they come in. The critical background story is that the optimal existence of classes of biological agents, and their individual members, entails the existence of a generative model (or generative models) in the agent's brain of the two-way causal interactions between its internal states (i.e., its neural economy) and the activities of its body and environment, e.g., how its internal states bring about causal changes to its body and environment, and vice versa. To best serve the homeostatic requirements of the agent, the generative model should be able to accurately predict the consequences of changes to its embodied states and its interactions with the environment before those changes are brought about, and constrain the activities and behaviours of the agent on that basis. Thus, the goal of the brain is to maximise its predictive powers to optimise the behaviours of the agent whilst respecting the morphological constraints set by its embodied configuration and this necessitates the minimisation, if not elimination, of prediction errors (mismatches between the anticipated changes that are set by the brain and the actual feedback). This principle applies across different timelines, e.g., evolutionary, developmental and real time and is the process through which we come into meaningful (perceptual and motoric) contact with the world around us. This model, I will argue, has far-reaching consequences for our concerns about how to build a theory of the general optimisation strategies the brain employs to optimise the outputs of perceptual and sensorimotor processes that will allow us to proceed towards an analysis of the individual and collective contributions that the body image and the body schema make to the optimisation of motor behaviour in the environment.

Chapter five will apply the insights provided by the prediction error minimisation framework to the concerns of chapter two pertaining to the conceptual analyses of the body schema. In this chapter I will argue that it is consistent with the opposing requirements from Gallagher and Vignemont for a conceptually satisfying analysis of the body schema. With

respect to Vignemont's model I will show how the interpretation of the body schema I offer provides a more plausible reading of her conception of the body schema in terms of action-oriented representation, and can make good headway on the important issue that her account raises nonetheless: the importance of prediction for optimising motor behaviours. The overarching conclusion is that with the background story about the implementation strategies of multisensory processing, my account can synthesise the important insights her model provides. With respect to Gallagher's account, I will claim that my interpretation is broadly consistent with his functional reading of the components of the body schema. I will then proceed to show how my interpretation meets additional requirements Gallagher sets for a theory of the body schema; how to explain the responsiveness of the body schema to the constraints placed on its operations by the individuals' embodied configuration, its environment and intentional purposes. The upshot is that it allows us to synthesise the insights Gallagher provides in an illuminating way. This line of argument is provided in anticipation of a line of objection that I expect Gallagher might level against my interpretation of the body schema in terms of a prediction-error minimising action-oriented representation. Gallagher is staunchly against the idea that mental and cognitive phenomena, in general let alone the body schema, should be defined in terms of mental representation. I will respond by claiming that whilst Gallagher's line of objection might be applicable to Vignemont's model, my account can withstand the criticisms.

Having provided an interpretation of the body schema in chapter 5 from the perspective of the prediction error minimising framework, I then apply the insights from this model to Gallagher's analysis of the body image as a basis for explaining a) how the multisensory processes of the body image bring us into meaningful perceptual contact with our body and b) how the body image and body schema work together to optimise behaviour. The objective is not just to make headway on Gallagher's model, but to use it as a framework for establishing a developmental theory of the interaction between the body image and the body schema with which we can begin to address two important deficits in developmental psychology pertaining to a) the lack of an analysis of the body image and the body schema and b) the need for an accompanying causal theory to explain how mental and cognitive phenomena develop and transform across different timelines.

The overarching conclusion is that the concurrent analyses of the body schema in cognitive and psychology opened potentially productive avenues of research pertaining to the optimal importance of multisensory integration for perceptual and motor processes. However, the direction in which the leads were taken has led to a dead end. Nonetheless, considering the

mounting interest in the optimisation strategies of perceptual and motor processing in Bayesian inspired cognitive science there is no better time to re-evaluate those discussions and I hope that bringing the theoretical resources these models provide to bear on Shaun Gallagher's analysis of the body schema, and the body image/schema distinction, gives it the proper acknowledgement it deserves.

Chapter One

Multisensory Integration, the Body Schema and Peripersonal Space: Recent Trends in Cognitive Psychology.

Introduction

The importance of synthesising sensory information from the multiple sense modalities, such as vision, audition, touch and proprioception, for determining the character of perceptual experience has been acknowledged throughout the history of philosophy and psychology. However, it is only in recent times that multisensory integration has been afforded its overdue status as a worthwhile research program on its own right. Research interests in contemporary multisensory integration were initially organised around a conceptual distinction between the *body image* and the *body schema*. To address the historic conceptual confusion over the concept of the body schema, the psychologists Nicholas Holmes & Charles Spence surveyed a wealth of empirical data which cuts across neurological studies, behavioural studies, and phenomenological studies about the multisensory underpinnings of body perception. I claim that the treatment of the concept of the body schema in this paradigm violates the intended definition of this concept and propagates the conceptual confusions associated with the distinction between the body schema and the body image. However, instead of acknowledging the possibility that the problems might be caused by their approach, Holmes & Spence cite the problematic nature of the body schema as the cause of the problem and conclude that the distinction should be abandoned in favour of an alternative model which unifies neuroconstructivism and Bayesian causal inference. Ultimately, I will conclude that the root of the problem for the approach to defining the body schema in cognitive psychology is threefold. First, the authors disseminate several readings of the body schema across their work that perpetuates the very problems they set out to resolve. Second, the call to abandon the distinction has fallen on deaf ears in current developmental psychology, and the misuse of this concept continues to propagate. Third, the alternative model they provide is promising but their account is a) under-developed and b) is best construed as a theoretical framework that the body image and schema can be incorporated into. This sets up the basis for chapter 6 of the thesis where I will show how the body schema and image distinction can be incorporated into their model.

1.1. Old and New Observations about the Multisensory Nature of Perception.

From our waking moment, we are continually bombarded with a plethora of sensory experiences, such as sights, smells, and noises, which are provided by the information from

our various sense modalities such as the visual, olfactory and auditory systems, respectively. As such, it would be uncontroversial to say that the perceptual experiences of humans, and other suitably constituted agents, are multisensory in nature. Consider the following example. As I swirl the spoon around the coffee that's in my cup I simultaneously see and feel my hand moving around in a circular motion and hear the noise of the spoon as it meets the inside of the cup. From my perspective, my visual, tactile and auditory experiences are not three experientially distinct events that just happen to occur simultaneously. Rather, they constitute the character of a unified experience of a particular event. The important questions that emerge pertain to *how* and *why* cross-modal inputs are integrated to produce a combined percept[1]. The former requires an explanation of the nature of the underlying operational principles that govern the interactions between the senses. The latter requires an explanation of the adaptive advantages, if any, that are brought about by the collaboration between the senses for perception. The former question has long entertained philosophers and psychologists who have sought to understand the nature of the process (or processes) through which the data provided by the distinct sense modalities produce a coherent percept.

Much ink has been spilt on the 'how' question over the history of philosophy. Proposals to this question can be traced back to at least Plato and Aristotle both of whom agree that perception requires not just the cooperation between the sense modalities in some loose sense, but rather the active integration between the distinct inputs for the purposes of producing a coherent percept that somehow combines the informational input from each. This type of claim is the working assumption of the multidisciplinary research program that is now known as 'multisensory integration', which seeks to understand the neural underpinnings and cognitive significance of cross-modal integration (more on this below). However, the views of Plato and Aristotle come apart over the issue of where sensory integration occurs. For Plato, the site of sensory integration is in the brain, whereas for Aristotle the input from the distinct modalities was linked by 'the common sense' (or intellect), a property he ascribed to the heart[2]. Perhaps the most influential precursor to modern ways of thinking about sensory integration comes from George Berkeley[3]. As part of his attempt at establishing a new theory of vision, Berkeley outlined various precursors that are now well established in the

---

[1] In this chapter I will concentrate on historic and current approaches to the 'how' question. Chapters 3 and 4 will provide an account of current approaches that attempt to answer the 'why' question.

[2] This is as much as I need to say about the Platonic and Aristotelian perspectives on sensory integration for the stage setting purposes of this section. For more, see P. Grigorov (1998) for a book length overview of the commonalities and differences between each school of thought.

[3] For more see Berkeley, 1709.

multisensory integration literature. The first is the idea of 'sensory capture' where the data from one sense modality is made to conform with the input from another in cases where the informational input that is provided by each modality is equivalent[4]. Second, as an empiricist, Berkeley proposes that sensory relations are experience dependent. This Berkeleyian model was influential for psychologists such as James Gibson[5] who, like Berkeley, tried to establish a new vision science which was based on the premise that the active interactions of sensory agents in their environment are driven by their information seeking sensory systems which extract invariant features of the environment. We will see in chapters 3 and 4 of the thesis that the Platonic and Berkeleyian proposals outlined here are implicit in the current theoretical models for understanding the dynamic nature of the multisensory processes that underpin perceptual and motor processes. First, there exists an overwhelming, and still growing, body of evidence which points towards the idea that a diverse range of mental and cognitive phenomena are dependent upon the synthesis of cross-modal data in various regions of the brain; second, it indicates that cross-modal information processing is responsive to the dynamical interactions between the body and the environment (more on this below)[6].

1.2 The Role of the Body Schema Concept in Contemporary Multisensory Integration Research.

The preceding considerations provide a rough sketch of the philosophical issues about the significance of multisensory integration that have echoed down through the ages. However, it is only in relatively recent times that multisensory integration has been afforded its status as a worthy research paradigm on its own right. Over the past three decades an outpouring of empirical data from studies on the superior colliculus in cats has provided important insights about where in the (feline) brain cross-modal integration takes place, and how cross-modal operations influence behaviour (e.g., Meredith, 1983, Meredith & Stein, 1993, & Calvert,

---

[4] More will be said about this in chapter 3 of the thesis.
[5] See Gibson, 1968.

[6] The reader should note that this assumption is by no means universally endorsed. Fodor claims that modality specific processes are encapsulated which means that they remain insulated from one another i.e., vision cannot affect tactile processing; and vice versa. This principle cuts across each of the sense modalities. This speaks directly against the possibility of sensory capture. However, contra Fodor, there is evidence on the table to suggest that cross-modal interaction happens at a very early stage in sensory cortical domains (e.g., Stein & Meredith, 1993, Stein & Stanford, 2008)

2004). More recently, this methodology has motivated further investigation into the significance of multisensory integration for a wide range of mental and cognitive capacities as they apply to (mature and developing) humans, primates, other non-human animals and robots. One key player who has influenced the general shape of more contemporary work in multisensory integration research in philosophy and psychology is the perceptual psychologist Charles Spence who has provided a plethora of useful surveys of empirical studies on the significance of cross-modal integration for body perception across a series of co-authored works (e.g., Holmes, Spence *et al,* 2003, Driver & Spence, 2004, Holmes & Spence, 2006)[7]. His work is of no small importance as any student or interested party who wishes to acquire a solid grounding in more contemporary issues in multisensory integration research will almost certainly need to familiarise themselves with the works of Charles Spence in one way or another (the relevance of this point is explained below). Nowadays, Spence focuses his attention on the general question of how the interactions between the distinct modalities are expressed in experiential consciousness, particularly how the causal interactions between visual, tactile, auditory and olfactory information influence or enhance taste perception, respectively[8]. However, the defining objective of his earlier work centred around two more specific questions about the significance of multisensory integration: 1) what role does the collaboration between the distinct modalities play in determining the adaptive motor responses of the body? 2) how is the perception *of the body* shaped by cross-modal interactions? To answer such questions, Spence and his colleagues bracketed them under a conceptual distinction between the *body schema* and *body image*, respectively.

Understanding the significance of this approach requires a brief explanation of the wider contextual motivations at play. At the time of writing, there was sufficient reason for thinking that questions about how cross-modal relationships shape behaviour and perception should be kept separate. For instance, the two-visual systems hypothesis from Goodale & Milner (e.g., Milner & Goodale, 1992) is grounded in empirical work that seems to indicate that the underlying visual processes of perceptual and motor processes are neurologically and functionally distinct. Therefore, to keep the questions Spence and his contemporaries were interested in apart, they categorised them under the body schema and body image distinction which is, in principle, supposed to provide a conceptual framework for understanding the sensory processes of body perception and motor control separate. Questions about the

---

[7] See Bremner, Lewkowicz, & Spence, 2012 for an edited collection on pre-existing and new advancements into the multisensory underpinnings of a wide range of perceptual, cognitive and behavioural capabilities, from balance and body perception in young infants to vocal communication in primates.
[8] See, e.g., Reinoso, Carvalho, Spence *et al.* (2016) & Spence (2015, 2014).

sensory underpinnings of motor control come under the concept of the body schema, whereas questions about the sensory underpinnings of body perception come under the body image. Around 2003 to 2006 Spence and his colleagues focus their attention on the first question and so their initial research is best understood as an empirical analysis of the body schema (I'll come back to this point). The concept of the body schema has been deployed across the cognitive sciences and beyond for quite some time after it was popularised just over a century ago by Henry Head & Gordon Holmes[9]. Head & Holmes refer to two distinct 'body schemes' that are 'organised models of ourselves'; so-called because they coordinate sensory information that is generated by the body for organising motor control and localising sensations on the surface of the body. First, there is the *body schema* which is understood to be a 'postural model of ourselves' that involves a 'combined standard against which all subsequent changes of posture are measured before they enter consciousness' where 'every new posture or movement is recorded on this plastic schema and the activity of the cortex brings every fresh group of sensations evoked by altered posture in relation with it' (Head & Holmes *et al.*, 1920, p.187). In other words, the body schema registers the sensory information that is generated by the postural changes of the body by comparing incoming information with the inputs from recent postural changes as the body proceeds through its current motor trajectory. Second, there is the 'superficial schema' that registers somatosensory inputs from tactile receptors on the surface of the skin. Importantly, the body schema and superficial schema operate *below the level of consciousness*. This means that their respective processes and outputs 'cannot be called into consciousness' (Head & Holmes *et al.*, 1920, p.187). Contemporary evidence for the existence of a non-conscious superficial schema can be found in empirical studies on a condition called 'numbsense'[10]. Individuals with this condition are unable to consciously feel tactile stimulation on their body, for instance they will not be able to experience the felt sensation of touching, or being touched by, an object. That being the case, they will be unable to feel someone touching the affected part of their body for the individual would have no qualitative experience of being touched. Nevertheless, experimental studies show that these individuals can accurately locate where they have been touched previously (in the absence of obvious visual feedback)[11]. This would suggest that tactile information is being registered, even if this information does not arise at the level of conscious experience. In turn, this speaks in favour of the existence of the

---

[9] Head & Holmes, 1911-12, Head *et al,* 1920.

[11] E.g., Rosetti, *et al.* 2001.

superficial schema. Similarly, the processes of the non-conscious body schema are revealed by the apparent fact that the body produces movements that do not require conscious deliberation or monitoring; indeed, some of those movements occur outwith the scope of conscious awareness. To validate this distinction Head & Holmes draw upon their studies on intramedullary lesions to the upper end of the spinal cord just prior to the junction between the dorsal column and the medulla oblongata. Such a lesion affects only one side of the spinal column. One such case is W.C. who, following a car collision, suffered such a lesion to the left side of his spinal column. As a result, he exhibited motor weaknesses to the right side of his body, but could nonetheless feel sensations like heat, pain and coldness on this side. On the other hand, the motor capabilities of the left side of his body remained preserved, yet this side of his body was insensitive to skin pricks, other tactile stimuli and heat. Broadly speaking, this leads Head & Holmes to conclude that the sensory underpinnings of the experience of sensations on the body must be neurologically and functionally distinct from those that facilitate motor control[12].

Importantly, if the postural schema can function independently of the sensory processes which underpin conscious sensations, and the superficial schema registers sensations below the threshold of consciousness, then there must exist a third system of sensory processes that can be 'called into consciousness' (Head *et al.*, 1920, p.187) to facilitate *experiential awareness* of the body. Clearly, we do have qualitative sensory experiences of our body. As we walk about, for instance, parts of our body routinely appear in our visual field and experience the felt sensation of its movement. Thus, we don't just have qualitative sensory experiences of what's out there in the world. Rather, some aspects of experience are directed the body. The question is what category such sensory experiences come under. They cannot be accommodated by the concept of the superficial schema for its outputs are not accessible to consciousness. They presumably cannot be accommodated for by the concept of the body schema either for the body schema pertains to the sensory processes that underpin the movement of the body, not the perception of it. Enter the concept of the body image to cater for this class of body experience (e.g., Lhermitte, J & Tchehrazi, E: 1937). Thus, sensations on the body can be registered yet do not arise to the level of awareness are categorised under the superficial schema; sensations which do are categorised under the body image. The important upshot is that a very broad distinction can be made between the sensory processes that are implicated in motor control (i.e., the body schema) and the sensory processes which

---

[12] See Head & Holmes (1911-12, 1920) for a more comprehensive overview.

are involved in the registration of sensations on the body, which may scale up to the personal level or not (i.e., the body image or superficial schema, respectively).

Nowadays, philosophers and cognitive scientists make use of the very general distinction between the body schema and the body image, and it is important to note that Spence and his colleagues follow suit[13]. For the provisional purposes of this chapter it suffices to say that the concept of the body schema should, in principle, refer to the consciously inaccessible system of proprioceptive processes that are involved in registering the position and movement of the body. The body image, by contrast, is the system of consciously accessible multisensory processes which shape and support personal level awareness of the body.

With these considerations in place, it is important to note that the distinction between the body schema and the body image can be drawn up in three distinct ways. The first concerns the sensory information each system works over. The body schema is a proprioceptive system, whereas the body image is a multisensory system. Secondly, whilst the processes of the body schema are inaccessible to consciousness, the processes of the body image are. Thirdly, whilst the body schema processes proprioceptive information for registering the postural changes of the body to appropriate motor control, the processes of the body image process multisensory information for structuring experience of sensations on the body. This much sums up the motivations behind, and the nature of, the body image and schema distinction[14]. As such, we can now proceed to consider the important issues for the distinction that emerge out of more contemporary work on multisensory integration.

1.2. The Body Image and the Body Schema: A Distinction without a Difference?

The preceding section lays out the general distinction between the body schema and the body image. This is important for understanding the rationale of Charles Spence and his co-authors who, just over a decade ago, inspired renewed interest in the concept of the body schema as an explanatory tool for organising multisensory integration research. They do so by surveying a wealth of research on the neural architectures in various cortical regions of the monkey brain that are simultaneously responsive to multisensory stimuli from the body and non-bodily objects located in the space that surrounds body. The spatial region the authors are

---

[13] We will see in the next chapter that the superficial schema is now considered to be part of the body image.
[14] In chapter 2 of the thesis we will consider more ways in which the distinction plays out in Shaun Gallagher's various analysis of the distinction.

interested in is peripersonal space which is the space surrounding the body that is within reaching distance of the effectors, relative to the current position of the body. The significance of this pertains to the fact that objects which exist in body space (i.e., the space occupied by the body) and peripersonal space are adaptive in ways that behaviours in relation to objects outwith these spatial regions are not. When objects are located outwith peripersonal space they are said to be in extrapersonal space. The important point is that non-bodily objects can be included as part of our body space when they are incorporated into our body, for instance when a fork is incorporated into the hand during mealtimes. Furthermore, objects that are in peripersonal space require a different class of behaviours compared with those in peripersonal space. For example, when, say, a coffee cup is in peripersonal space an individual can retrieve it by reaching out their hand and arm, but if it is in extrapersonal space they may have to get up and walk towards the location where the mug is, such as the coffee table. Furthermore, objects in peripersonal space have a different adaptive significance to objects located in extrapersonal space for they can potentially present a bigger threat. Objects close to the body present a greater risk for hazardous collision than objects which are far away. Therefore, the authors speculate that the responsiveness of the brain to objects in peripersonal and extrapersonal space must be different. Furthermore, they claim that accommodating for the on-going changes to the spatial relations that obtain between the body and non-bodily objects requires a body schema that is regularly updated to facilitate appropriate behavioural responses in relation to objects when they are in body space, peripersonal and extrapersonal space. As such, Spence and his colleagues seek to answer the question of whether the brain is responsive to objects in peripersonal space in a way that it is not towards objects in extrapersonal space via an empirical analysis of the neural underpinnings of the body schema. This approach has an additional advantage for an investigation into where the body schema is in the brain should help to resolve age-old problems that had surrounded the body schema since its conception. At the time of writing, the distinction between the body image and body schema had been falling in and out of favour across the cognitive sciences for some time because the persistent and pervasive conflation between the two concepts caused deep seated terminological and conceptual confusions. Such problems had been acknowledged for some time[15], yet remained unaddressed. Here is a particularly striking example:

---

[15] E.g., Paillard *et al.* (1972), Gallagher, (1992, 1998, 2005)

> '*The image of the human body* means the *picture of our own body* which we form in our mind, that is to say the way the *body appears to ourselves…We call it a schema of our body or bodily schema,* or following Head, …*postural model* of the body. *The body schema is the tri-dimensional image everybody has about himself.* We may call it the 'body image' (Schilder, 1935, p. 11, italics for emphasis).[16]

This is a very clear example where the body schema is simultaneously described in terms of an 'image' or 'picture' of the 'appearance' of the body or 'oneself' that's formed in the mind, or indeed as a kind of 'body image'. Clearly, this only serves to collapse the distinction[17]. However, an equally problematic implication is that because the body schema is so ill-defined it is not easy to mount a clear research project to investigate the nature of the body schema. To make headway with respect to understanding what the body schema is and how it functions the authors propose their own definition of the body schema (more on this below) and mount a clear research project which allows them to pursue a clearly defined investigation into its nature. I'll explain how they attempt each goal in the ensuing sections.

1.3. Evidence for the Existence of the Body Schema from Single-Cell Studies in Monkey Cortex.

With the intention of remaining faithful to the definition from Head & Holmes, Spence *et al.* 2003 start with a provisional definition of the body schema in terms of an 'integrated multisensory representation of the body' in the brain. So defined, the body schema is reducible to the neural architectures that are responsive to cross-modal stimuli from the body. As such, the distinctive hallmark of their definition is that, contra Head & Holmes, the processes of the body schema are multisensory, not proprioceptive.

The important upshot is that a comprehensive story about what the body schema is and how it operates requires an account of not just the neural populations that are responsive to stimuli from the body, but also how incoming streams of sensory relations calibrate and re-calibrate motor behaviour in the kinds of ways the preceding examples suggest. To address this very issue Holmes & Spence[18] summarise a comprehensive body of empirical data from single cell studies in monkeys. The first class of studies investigate the responsiveness of the

---

[16] I borrow this example from Shaun Gallagher's *2005*: p.20.
[17] For a more comprehensive overview see ch.1 of Gallagher's *2005*.
[18] Holmes, Spence *et al.,* 2003.

neurons in various brain regions to stimuli from the body (see below) which, the authors suggest, is important for appropriating behavioural responses in the environment. For instance, a class of neurons in the cortical regions of the monkey brain which are thought to be involved in encoding the posture and movement of the body and the planning and execution of movement, such as Brodmann's areas 5 and 2, were discovered to be simultaneously responsive to visual and proprioceptive information about limb position (Graziano *et al*, 1999). Because the neurons are responsive to information from at least two modalities they are referred to as 'bi-modal' neurons. Bimodal neurons, which are responsive to tactile input from the head and auditory input immediately surrounding the head, can also be found in the ventral pre-motor cortex (Graziano *et al* 2002)[19]. Similar studies (e.g., from Graziano *et al* 1995) investigated the responsiveness of visuo-tactile neurons in several interconnected brain areas including inferior area s6 in the frontal lobe and 7b in the parietal lobe, as well as the putamen in the monkey brain. The receptive field of the bi-modal neural circuits in question is responsive to tactile stimulus that is produced by stroking the hair located on the arm, hand and face and visual stimuli surrounding each respective body part up to twenty centimetres beyond the skin of the arm and hand and ten centimetres beyond the face (Graziano, M.S & Gross, C, 1995, p.1023). Therefore, while the tactile receptive field of these neural circuits is limited to the surface of the skin, their responsiveness to visual stimuli extends beyond the body. Furthermore, the responsiveness of the neurons to visuo-tactile stimuli is body-part specific. For example, the neurons that are responsive to tactile stimuli from, say, the hands are only responsive to visual stimuli from within twenty centimetres of these parts of the body and no other. This is demonstrated by an increase in the activation rate of the neural circuits that are responsive to tactile stimuli from a specific body part whenever visual objects move towards that body part. Likewise, their firing rate decreases as the object moves outwith their visual receptive field (Graziano, M.S &Gross, C, 1995, pp.1023-1025). Furthermore, the visual field of body-part specific classes of neurons shifts relative to the movement of the relevant effector. For example, their firing rate increases and decreases when the relevant body part is brought closer to, or farther away from, a visual object (Graziano, M.S & Gross, C, 1995, p. 1024). The crucial conclusion is that bi-modal neurons are simultaneously responsive to stimuli from the body *and* objects in peripersonal space, but not extrapersonal space. The only way these neurons can become sensitive to an object that is currently located in extrapersonal space is by bringing the object into peripersonal space.

---

[19] See Holmes & Spence, 2004, for a more comprehensive overview.

After summarising these kinds of studies, Holmes & Spence draw several important conclusions about where the body schema is distributed in the monkey brain, such as Brodmann's areas five and two; the parietal lobe; the ventral premotor cortex, and the putamen, respectively. Furthermore, the fact that the neural populations in these regions are bi-modal, and in some cases tri-modal, means they are multisensory neural populations. Furthermore, these neural populations are responsive to stimuli from the body (i.e., body space) and peripersonal space, but not extrapersonal space. Thus, if the neural regions just specified can be understood as part of the monkey body schema we can say that it can be understood in terms of the multisensory neural architectures that are sensitive to the body and objects within proximity to the body.

1.4. Evidence for the Functional Plasticity of the Body Schema: From Monkey Neurons to Rubber Hands.

The next step Holmes & Spence take is to show that the body schema exhibits a functional plasticity of sorts in the sense that the bi-modal properties of the body schema are flexibly responsive to stimuli from objects that are within reach of a tool when it is incorporated into the body when it is out of reach of the effectors. In the preceding section I gestured towards the idea that incorporating a tool alters the spatial dimensions of our body. We can use tools to alter the spatial dimensions of our body in various ways. For example, incorporating a tool into our hand lengthens our arm and extends its reaching distance. If my ball falls into a river and I can't reach it with my hand without the risk of falling into the water, I may use a stick instead to try to retrieve it. Thus, the guiding question which underpins the review of the second set of studies (more on this below) asks whether the neural mechanisms that are responsive to visual objects that are within reach of the effectors readapt to non-bodily objects when are temporarily incorporated into the body for task specific purposes. The short answer is 'yes'. In support of this claim Holmes & Spence review studies from Atsushi, Iriki, Tamaraya et al[20] that investigate the sensitivity of bimodal neurons (visual and tactile) in the intraparietal area of the cerebral cortex in the brain of macaque monkeys to objects within the peripersonal spatial region of the effectors and small hand-held rakes (e.g., Iriki et al. 1996).

---

[20] Iriki et al. 1996

At the first stage of the experiment, the responsiveness of the bi-modal neurons to visual stimuli from objects within the peripersonal spatial region of the effectors was tested for by placing a raisin within reaching distance of the monkey's hand. In a similar vein to the studies from Graziano *et al.,* it was found that the responsiveness of the neurons to the presence of the raisin increased and decreased as the raisin was placed closer to, and farther away from, the hand, respectively. At the second stage of the experiment, the raisin was placed outwith the reach of the monkey's hand and, as would be expected, this was met with a radical decrease in the activation rate of the neurons and in some cases they stopped firing completely. At the third stage of the experiment, the monkeys were trained to use a small rake to bring raisins that are outwith the reach of their hand. The important upshot is that when the raisins are placed outwith the maximal reaching range of their hand the neurons would show little to no activity. However, after being trained to use the rakes to bring the raisin into reach, the activity of the neurons would reactivate when the raisin was within reach of the rake. What this shows, so the story goes, is that the cross-modal relations that underpin the localisation of the body and objects in peripersonal exhibit a sort of functional plasticity insofar as the visual responsiveness of the neurons can readapt to tools when they are incorporated into the body to fulfil a given task. I use the term 'functional plasticity' to emphasise that the readaptations of the underpinning cross-modal relations is purely functional in the sense that inputs from each modality can be coupled and de-coupled for the requirements of a specific task. This is to distinguish the kind of plasticity in mind from the more traditional understanding of neuroplasticity in terms of the more substantial, and sometimes permanent, re-routing of the material architecture of the brain. The important conclusion that Holmes & Spence draw from studies of this kind is that if the same multisensory neural circuits which are responsive to stimuli from the body and objects in surrounding space are classed as part of the body schema of macaque monkeys, and those neural regions also readapt to tools when they are incorporated into the body for some task specific purpose, then the body schema exhibits functional plasticity.

The preliminary conclusions Holmes & Spence draw from the single-cell studies are that the neurological substrates of the body schema in the brain of macaque monkeys, such as the inferior area 6 in frontal lobe, area 7b in parietal lobe and the putamen, are multisensory; responsive to the position of the body and objects in peripersonal space and exhibit functional plasticity.

Despite the important insights that the studies provide into the nature of the body schema in monkeys, the reader might still question how we can make similar observations about the

human body schema. To address this, Holmes and Spence outline a study that would seem to demonstrate that the same kinds of re-adaptations to non-bodily objects that can be observed at the neurological level in the studies on macaque monkeys could also be observed at the experiential level in humans. The study they appeal to is the now well-known rubber hand illusion (e.g., Botnivik, 1998). During this experiment one of the participant's hands is placed behind a screen so that he or she can no longer see their hand. A realistic looking, human-like, rubber hand is then situated in its place. Afterwards, the experimenters introduce identical tactile stimulation to the participant's occluded hand and the rubber hand. Therefore, the subject *sees* the rubber hand being stroked but *feels* the sensation on their real hand that is out of their view. The important result from the experiment is that when the subject is asked to locate the position of the tactile sensation, the participant will indicate that they feel the sensation on the rubber hand. In effect, the rubber hand illusion creates competition between visual and tactile input about the spatial location of the stimulus, as well as a disconnection between visual and proprioceptive input about the position of the hand. The rubber-hand illusion provides a clear illustration of the phenomenon of sensory capture that was mentioned earlier. In this case, tactile and proprioceptive inputs are made to conform to the visual input about the location of the tactile stimulus and the position of the hand, respectively[21].

What this gives us, Holmes & Spence claim, is an insight into the different ways in which the body schema is responsive to non-bodily objects. The first set of single cell studies demonstrates the multisensory and functionally adaptive nature of the body schema at the neurological level, and the rubber-hand illusion reveals similar properties of the human body schema (Holmes & Spence, 2004, p. 101).

---

[21] See F. De Vignemont (2009, 2010) and Carruthers, G (2008, 2009) for a critical discussion. Vignemont questions the reliance on retrospective verbal reports in such cases. Carruthers invokes a distinction between having a sense of embodiment (the sense of the boundaries of one's body) and the sense of body ownership (the sense that a body part belongs to one's own body) which is important when it comes to evaluating the implications of the rubber hand illusion that I must skip over here. However, I will challenge the assumption that interpreting this kind of study provides insight into the body schema below.

1.7. Problems for the Interpretations of the Body Schema from the Empirical Studies of the Multisensory and Functionally Adaptive Properties of the Cortical Regions in Monkeys.

The preceding considerations make it clear that the approach to explaining the nature and functional properties of the body schema perceptual psychologists take largely involves identifying the neural populations that are responsive to cross-stimuli from the body and objects within reach of the body and analysing some of their properties, such as their responsiveness to objects that are incorporated into the body. Such an approach not only allows us to pin down (some) of the underlying neural architecture of the body schema, it also allows us to say something about its nature (as specified above). The explanatory goal of this approach is to better understand the multisensory mechanisms that are responsive to stimuli from the body and non-bodily objects (i.e., as they approach or are incorporated into the body). As such, Holmes & Spence claim that this approach takes cognitive scientists some way towards addressing perennial questions about where the body schema is in the brain and how it operates.

However, the implications of this kind of approach are not as straightforward as Holmes & Spence suggest, and they acknowledge this point in their *2004*. To understand the significance of the ensuing commentary, remember that the principled basis on which the distinction between the body schema and the body image is drawn is the fact that the processes of the body schema are *not* accessible to consciousness, whereas, the processes of the body image can, in principle, be consciously accessed. Holmes & Spence recognise the possibility that visual *attention* and the *expectation* of reward might play a significant role in determining the results of the single-cell studies. For instance, in the first set of studies on the responsiveness of bimodal neurons to tactile stimuli from the body and visual stimuli from nearby objects, the experimental design meant that the monkeys' visual attention was drawn to a particular body part and a visual stimulus (a cotton bud in this case). In the second set of studies it is not clear whether the excitability of the neurons was generated by the visual stimulus or the expectation of a treat. The possible influence of personal level capacities like attention and expectation of rewards makes it difficult to make decisive interpretations from the data in relation to the body schema that are supposed to work independently of personal level processes[22]. I would also add that the rubber-hand illusion comes up against the same

---

[22] This claim will be qualified in the next chapter. Shaun Gallagher's analysis of the distinction demonstrates that whilst this assumption is necessary for making a conceptual dissociation between the body image and the

problem for the experiment gives us access to the results of the readaptations to the cross-modal relationships which underpin the individual's sense of where an individual experiences tactile sensations on his or her hand, and their sense of where their hand is located. As such, the rubber hand illusion experiments might be best considered as not so much an illusion that is generated by the readaptations of the multisensory processes of the body schema, but of the body image. The question, then, is how perceptual psychologists should proceed with their investigation into the body schema if the data doesn't quite fit their conception of the body schema. I turn to the various approaches Holmes, Spence and their colleagues take in the next section.

1.8. Problems for the Working Concept of the Body Schema in Perceptual Psychology.

From the fact that the empirical evidence they survey is apt to cause confusion between the body image and the body schema, Holmes & Spence (and their collaborators) eventually come to the conclusion that rather than tackle the 'slippery issue' of clearing up the conceptual confusion between the two concepts (Holmes & Spence: 2006, p.6), the distinction between the body image and the body schema should be abandoned from the multisensory integration theorist's conceptual toolbox altogether (e.g., Holmes & Spence: 2004, 2006). Instead the multisensory integration theorist should seek an alternative explanatory framework. The proposed alternative will be outlined below. First, I want to address their first claim about the future of the distinction in cognitive scientific theorising.

The reader may have some degree of sympathy with this suggestion because the distinction doesn't lend itself easily to a clear interpretation of the aforementioned empirical data. However, I propose that the grounds upon which Holmes & Spence make their claim are problematic for several reasons. Their reasons for rejecting the distinction based on age-old problems that are associated with the body schema are unwarranted because by the time the authors get around to this suggestion, they had already disseminated various definitions of the body schema (see below). This is no small matter for in section 1.3 of this chapter I mentioned the importance of their work in relation to the body schema for getting to grips with contemporary research interests in the multisensory integration research. Any student, of

body schema, the two systems interact during the course of everyday motor and perceptual processing. However, this acknowledgement doesn't undermine the general thrust of the current line of argument.

whom it is required to have a basic grounding in contemporary multisensory integration research, will be introduced to the key research paradigms in terms of an empirical analysis of the body schema. However, no two definitions of the body schema are ever the same. As such, I claim that they cannot appeal to old problems in support of the rejection of the concept, and the concept of the body image with it when they are very much responsible for propelling the same confusion forward into the most contemporary discussions about the body schema at the time. Let me explain. In their seminal *2003* Holmes *et al.* offer the following definition of the body schema:

> The body schema has often been invoked as an explanatory concept when it should perhaps rather be considered as a *label for a set of problems* still requiring explanation. (Holmes, Marivata, & Spence, 2003, p. 1, emphasis added)

The kinds of problems in mind are as-yet unanswered questions about how the causal relationships between cross-modal inputs support and transform body perception and motor control in adaptively important ways, in which case there is no need for a distinction between the body image and the body schema as questions that are typically segregated under these categories can be bracketed under the term 'body schema'. However, the empirical research we have just considered from their *2004* is framed around an understanding of the body schema in terms of an 'integrated multisensory representation' in the brain that is responsive to the body and objects that are within proximity of the body.

What Holmes & Spence fail to recognise is that it is the inconsistent definitions of the body schema they work with that leads them to their problematic conclusions, not the historic conceptual confusions that have always surrounded the distinction. Furthermore, the proposal to jettison the distinction from the literature has done nothing to quell the use, nor the misuse, of the concept of the body schema. For instance, Mendoza defines the body schema as the '*personal awareness* of one's body'[23]. Likewise, Reid defines the body schema as in terms of an '*understanding* of the biomechanical constraints inherent in the physical makeup of the human body'[24]. In doing so, they attribute properties to the body schema that are usually the province of the personal-level perceptual and cognitive processes that are, typically,

---

[23] Mendoza: 2011, p.427.
[24] Reid, p.153 in Virginia & Slaughter (eds.): 2014.

attributed to the body image. The take home message is that conveniently sidestepping the hard task of cleaning up the conceptual confusion by disposing of the distinction does not magically make the concept of the body schema and its associated problems go away. The issue is how perceptual psychologists should proceed from here, and it seems that they have two choices: they can either 1) make a more concerted effort to utilise the distinction correctly or 2) provide a viable alternative model that can do the same kind of work as the body image and schema distinction. We will see in chapter six of the thesis that developmental psychologists propose the first option. However, to my mind it seems that Holmes & Spence aim for the second. In the remainder of this chapter I will outline their proposal and (briefly) assess whether their model is a natural competitor to the body image and body schema framework.

1.9. A Neuroconstructivist Account of the Multisensory Underpinnings of Body Perception: A Satisfying Alternative to the Body Image and Schema Distinction?

In their *2006* Holmes & Spence suggest that the best way to move forward is to replace the body schema and body image model with an alternative classification system that categorises questions about the causal significance of cross-modal processing for body perception and motor control (Holmes & Spence, 2006, p.3). In their *2012*[25], Holmes & Spence give some indication of the kind of theoretical model that could be interpreted as an alternative to the body image and schema model: Neuroconstructivism (e.g., Mareschal *et al*, 2007). Neuroconstructivism is not an explanatory theory. Rather, it is a prescriptive model for what an explanatorily complete model of mental and cognitive phenomena should look like. Neuroconstructivists seek to explain how the structural and functional properties of the brain are developed, shaped and transformed by the interactions between the embodied states of the agent at all levels of description, e.g., the interactions between individual neighbouring neurons, domain-general and domain-specific brain regions, as well as the interactions between the brain and the body, and the interactions of the body in the (physical and social) environment. Importantly, the explanatory aim of neuroconstructivism is to synthesise the insights from empirical studies that investigate how body perception and motor control is

---

[25] Bremner Spence & Lewkowicz (eds.): 2012.

enabled and constrained at different levels of explanation, e.g., genetics; cellular, neurological and behavioural studies; embodied cognition; phenomenology and social cognition. Furthermore, it specifies that a complete explanation of a particular mental or cognitive phenomenon, and mental and cognitive phenomena in general, requires a supplementary story about the causal processes that drive the interactions at the specified levels of description across different timelines (e.g., in real time, and across evolutionary and developmental timeframes). Thus, a story about the nature of a given mental or cognitive phenomenon must tell a story about how its emergence at one level of explanation (e.g., at the behavioural level) arises out of the activities of processes at another (e.g., the neurological). Furthermore, the underlying causal story about how the explanandum emerges as it does must explain how this process is developed and transformed at different stages of development (I return to this point in chapter 6). Clearly, neuroconstructivism isn't a theory of what those causal processes might be for it only specifies that an explanatory model *should* have something positive to say about the nature of those processes and how they operate.

In their *2012*[26] Bremner, Holmes & Spence adopt a neuroconstructivist approach to explain the transformative effects that conflicts between cross-modal relations have on the developing infants' ability to (visually) localise tactile sensations on moving parts of their body (more on this below). Therein, they also give some indication of what they take to be the most promising supplementary story about the underlying causal processes at play: Bayesian inference e.g., Bernardo & Smith, 1994; Gelman *et al*, 2004). Bayesian causal inference is a mathematically specified process by which the brain actively anticipates sensory signals before they are delivered by the peripheral sensory systems in an on-going way. To do so the brain brings relevant prior-knowledge about the most probable cause of the current stimulus based on the nature of the sensory data and any relevant contextual information (more on this in chapter 3). Therefore, under this framework the nature of sensory processing is determined by feed-forward, probabilistic prediction.

By supplementing neuroconstructivism with Bayesian causal inference, Bremner, Holmes & Spence attempt to tell a story about how the development of the capacity to experience sensations on moving parts of the body is underpinned by Bayesian style expectations about the effects of bodily movement on visuo-tactile relations. To explain the underlying rationale of the approach the authors take up consider the following scenario. When a particular limb is in its typical position it normally occupies the corresponding spatial region of the field. For

---

[26] Holmes, Spence & Lewkowicz: 2012.

example, the right hand normally occupies the right side of the visual field, and vice versa for the left hand. Therefore, there is a direct spatial correspondence between the spatial position of a body part and the spatial region it occupies in visual space. As such, an individual would expect that tactile stimulus from a specific body part can be visually accessed at the corresponding region of his or her visual space. For example, suppose that his or her left hand is in its typical resting position (i.e., at the left side of their body), then they will understand that visual contact with an object that is felt pressing against, say, the left hand can be brought about by glancing to the left. Of course, there is an exception to this rule when an individual crosses his or her body part over to the opposite side such that it now occupies the opposite side of visual space. For example, when an individual crosses their left hand over to their right, it now occupies the part of their visual space of their visual space that is normally occupied by the right side. In this instance looking to the right, not the left, brings about visual contact with the left hand.

As mundane as this line of thought may seem, the studies from Bremner & co[27] demonstrate that whilst human adults are capable of flexibly readjusting their visual responses to the movement of their body, this ability only develops in young infants at some point between their sixth and tenth month. The authors use the preferential looking paradigm to investigate the developmental stages at which this adult-like ability to visually track tactile sensations when their body parts are spatially displaced is absent and present in young infants. The short version of the outcome is that this capacity is absent in infants at six-months-of-age and present in young infants aged 10-month. This is demonstrated by the fact that once the left hand is crossed over to the right the younger infants glance to the left when a tactile stimulus is introduced to their left hand, whereas the older infants correctly look to their right. The explanation Bremner, Holmes & Spence provide is that the visuo-motor responses of the younger infant are shaped by expectation about how visual access to their left hand can be brought about which has been shaped by previous multisensory experiences when their left hand is in its typical position. In this case, there is no conflict between vision, touch and proprioception about the location of their left hand. However, when the left hand is crossed over and introduced to some sort of tactile stimulation this generates new tactile-proprioceptive relations which requires a novel visuo-motor response that is sensitive to the change of position (Bremner, Holmes & Spence, 2012, p.120) Therefore, the visuo-motor response of infants at six-and-a-half-months exhibit an insensitivity to changes to the position

---

[27] Bremner, Spence *et al*: 2008.

of the hand. Nevertheless, sometime between the age of six-and-a-half and ten-months-of age the developing infants progress beyond their reliance on multisensory information about the limbs when they are in their typical position and exhibit visuo-motor responses which demonstrate a sensitivity to the movement of their body. The different visuo-motor responses at these stages of development, the authors speculate, is driven by the experience of changes to the systematic causal relationship between the displacement of the body parts and their visuo-tactile-proprioceptive experiences of the changing positions of their body parts. For instance, the six-month-old infant still *expects* to see their left hand in its typical position, even when it is crossed over to the right, however, infants aged ten-months expect that tactile stimulation from the left hand can be visually accessed by looking to the region of visual space it occupies when it is placed to their right. In turn, this expectation is determined by the belief that changes to proprioceptive feedback from the body which reports a spatial displacement of the limb to the opposite side of the body will drive a change in visuo-tactile expectations about the seen and felt location of the stimulus. Therefore, the different visuo-motor responses of the two age groups is anchored in an experienced based belief about the causal effects of body movement on sensory stimulation and visuo-tactile-proprioceptive experience. The presence or absence of this belief determines the prior expectation that a change in body position requires the same visuo-motor response that would be generated if the body part were in its typical position; or a different visuo-motor response to locate tactile sensations on the displaced body part. The behavioural responses of the older infants would suggest that they have this belief while the younger infant is yet to acquire it[28].

The overarching point is that the experimental paradigm is neuroconstructivist in nature because the experimental designs provides access to the results of the interplay between the sensory modalities via observations of visuo-motor behaviours, and a Bayesian inspired story about the underlying causes of the observed behavioural differences in infants at different stages of development allows the authors to provide a plausible story about one of the processes which drive the capacity to visually track tactile stimulus on the moving body.

The question is whether this model provides a satisfying alternative to the body image and body schema distinction. In one sense, the answer is 'yes' for it holds promise of providing a story about the underlying, motor-generated, causal processes which shape and transform the multisensory processes that determine perceptual apprehension of the body. In this sense, a neuroconstructivist model provides a rough-and-ready guideline for an approach to

---

[28] We will see in the succeeding chapters that Bayesian inference need not be conscious. Therefore, talk about the prior belief or expectation of the agent can just mean prior information that it stored in its brain.

understanding the underlying processes of the body image. However, in another sense the answer is 'no' for this one example hardly constitutes a full fleshed theory of what a neuroconstructivist/Bayesian model should offer in the place of the body image and body schema model insofar as there is, yet, no satisfying interpretation of the distinction on the table. At most, this model provides a brief snapshot into the kind of approach that might be useful for analysing the developing body image. However, this doesn't quite take us as far as explaining how the processes of the body schema are shaped and transformed, nor its relationship with the body image, and vice versa. The question of whether this model qualifies as a competing model to the distinction is still very much open to question.

However, there are reasons for thinking that instead of conceiving the body schema/image distinction and the neuroconstructivist/Bayesian models as opposing frameworks, the former can act as an important explanatory tool for the purposes of the latter. As we will see in chapter 6, developmental psychologists are looking to the body image and schema distinction to do the kind of explanatory work Holmes & Spence aim for with their neuroconstructivist/Bayesian paradigm. The central issue concerning developmental psychologists now is that while much ink has been spilled on the issue of how to dissociate the body image and the body schema, there is a dearth of theory concerning how the body image and body schema interact. Such a story is required, so the story goes, to understand scale errors. Scale errors are motor errors that infants make because of their failure to apprehend the scale of their body, such as its size. For example, they may attempt to enter a car that is the same size of their foot.[29] The lack of theory about the relationship between the body image and the body schema, and the lack of a causal story about how the interactions between each system work to shape and influence perceptual and motor processes, is an important deficit in developmental psychology. My first step towards addressing these issues will involve revisiting the philosophical analyses of the body image and schema distinction.

Conclusion

In summary, Holmes & Spence, along with their co-authors, have played a key role in bringing the concept of the body schema back into prominence in perceptual psychology as part of their attempt to organise and interpret empirical data about the multisensory nature of

---

[29] I develop this line of thought further in chapter 6.

the neural regions involved in responding to stimuli from the body and peripersonal space in monkey; behavioural studies which analyse the systematic causal relationship between bodily movement, the location of sensory stimulation and perceptual experience (i.e., the rubber-hand illusion), and developmental studies which investigate the developmental processes which underpin the capacity to visually track tactile stimulus on the static and moving body. However, their handling of the concept of the body schema only serves to further perpetuate the kind of conceptual confusion that already surrounded the concept (i.e., the confusion between this concept and the concept of the body image). By failing to recognise this, Holmes & Spence wrongly attribute the problems their approach fosters to the body schema and insist that the distinction between the body schema and body image is dropped in favour of a neuroconstructivist model of development that is supplemented by Bayesian causal inference. Up to now, their proposal for the distinction to be retired has gone unchecked and I have argued that it is unwarranted. Furthermore, their supposed alternative model remains underdeveloped. At this point in time residual confusions over how to understand the body schema concept, and its relationship with the body image, remain. Furthermore, I indicated that developmental psychologists have called out for a proper analysis of the distinction. As such, there is no better time than now to revisit the philosophical analyses and arguments about how the distinction should be drawn up. This will be the topic of the next chapter.

Philosophical Analyses of the Body Image and Schema Distinction

Introduction

The distinctive hallmark of philosophical analyses of the body schema is their inclusion of a motor component that is directly involved in the organisation and production of motor output. As such, the body schema doesn't just plot the spatial layout of the body; it is also directly involved in producing and regulating behaviours. Another key emphasis is the role of the body schema in *optimising* (human) behaviours. This issue emerges in Shaun Gallagher's influential analysis of the body image and body schema distinction, as well as the philosophical debates his account has inspired about how to conceptualise the body schema in a way that will deepen our understanding of its nature and its functional role. In this chapter I will rehearse Gallagher's analysis of the distinction and outline a critical response from Frederique de Vignemont in relation to his interpretation of the body schema. I will conclude by drawing out the important questions that the respective interpretations of the body schema open. I will conclude that Vignemont's critical response to Gallagher is baseless. On Gallagher's behalf, I will argue that her interpretation of the body schema does not possess the additional explanatory power Vignemont thinks it has. I finish by showing how her account of the body schema in a later paper only serves to further perpetuate the conceptual confusions that surround the body schema. The overarching conclusion is that the direction in which Vignemont has taken the existing philosophical analyses of the body schema has left potentially fruitful avenues of research unexplored and the confusions over the body image and body schema distinction unresolved.

2.1 The Objectives and Implications of Shaun Gallagher's Analysis of the Body Image and Schema Distinction: A Brief Summary.

From the considerations that were discussed in chapter one, it seems clear that if we are to avoid the confusions that are traditionally associated with the concept of the body schema then we must go back to basics and refer to the intended definition from Head & Holmes. So defined, the body schema is a system of sensory processes that are responsible for registering the postural configuration of the body, and for comparing new inputs which are generated by the changing postural disposition of the body during movement. Somehow, though their account doesn't quite extend so far as to explain how, these processes of the body schema play an important role in organising and facilitating motor output *below the threshold of personal-level awareness*.

In response to the same kinds of problems that motivated the approaches to redefining the body schema in perceptual psychology, and to build on the story told by Head & Holmes, the philosopher Shaun Gallagher has been attempting to clear up the conceptual confusions that have plagued the distinction since at least the 1980s[30]. Broadly speaking for now, he does so by attempting to provide a clear definition of the body schema that remains faithful to the definition provided by Head & Holmes that is rooted in principled lines of demarcation between the body schema and body image. Crucially, what sets Gallagher's analysis apart from the approach in perceptual psychology is that once he makes the conceptual dissociation between each system, he then proceeds to provide an analysis of the relationship between the body image and body schema during every day perceptual activity and motor behaviour. In turn, this analysis provides a revealing insight into the individual and collective contribution that the body image and the body schema make to the optimisation of perception and action. I will return to these important strands of thought at various stages throughout this chapter. First, I turn to Gallagher's interpretation of the body schema (as a concept); his account of what the body schema (as a system) consists of, and how it should be differentiated from the body image.

2.2 Gallagher's Functional Interpretation of the Components of the Body Schema

For ease of exposition, I will mostly concern myself with the account of the human body schema and analysis of the body image and schema distinction Gallagher provides in his *2005*. To my understanding this is the latest, and most up-to-date, version of his framework. In accordance with the account from Head & Holmes, Gallagher concedes that the behaviours produced by the body schema are *prenoetic* by which he means the operations of the body schema do not require personal-level intervention and/or monitoring (see, e.g., Gallagher, 2005, pp.24-26). In the interests of clarity consider the following examples from Gallagher. Whilst you reach over to pick up your coffee cup you may be visually and proprioceptively aware of the movement of your arm as it extends outwards toward the cup. In other words, you may simultaneously see and feel your arm moving towards the cup. Meanwhile, you may be unaware of your fingers curling around to form a grasp in the precise way that's required

---

[30] E.g., Gallagher: 1986, 2005 & Coles & Gallagher: 1998.

for picking up the cup. This motor adjustment is carried out regardless of whether you are aware of it. Likewise, you may rapidly duck your head out of the way of an overhanging tree branch whilst your attentional focus is directed elsewhere, such as the engrossing conversation you are having with your friend as you walk along the pavement. On reflection, such motor responses seem to be not so much volitional as automatic. However, Gallagher is clear about the fact that the behaviours produced by the body schema are *not* automatic. Retrospectively, it may seem as though they are automatic in the same way reflex movements are, but this would be to underestimate the complex neurophysiological underpinnings of those motor behaviours. Rather, the behaviours produced by the body schema are mediated by a complex matrix of top-down influences (such as intentions) and multisensory calibration, as well as bottom up influences such as the environment (more on this in chapter 5). For now, it suffices to say that for Gallagher, integrated sensorimotor processes underpin the motor behaviours that are produced by the body schema. To explain Gallagher characterises the body schema in functional terms. In other words, Gallagher provides a description of the entire body schema system, and its sub-components, in terms of their function. For instance, here he describes the body schema in terms of:

> […] *a system of motor capabilities, abilities and habits that enable posture and movement.* The body schema is not a perception, attitude or belief. *Rather, it is a system of motor and postural functions that operate below the level of self-referential intentionality, although such functions can enter into and support such intentional activity. The preconscious sub-personal processes of the body schema system are tacitly keyed into the environment and play a key role in governing posture and movement.* Although the body schema can have specific effects on cognitive experience…it does not have the status of a conscious representation or belief. (Cole & Gallagher, 1998, p.372, emphasis mine).

The italicised parts of the passage are the aspects of Gallagher's analysis of what the body schema is that I will be concerned with at this stage[31]. The body schema is a system of motor 'capabilities', 'habits' 'functions' and 'performances' of the body which enable the maintenance of balance and movement. What, then, does this mean? With respect to the body schema, perhaps it would be more accurate to say that the body schema is a system of motor functions that underpin our motor capabilities, habits, and performances. For in his *2005*

---

[31] Gallagher's commentary on the role that the body schema plays in supporting intentional activity and cognitive experience will be explained in more depth in chapter 6.

Gallagher probes more deeply into the details of what he means by a behavioural capability, habit and performance which he explains in terms of a pattern of sensorimotor processes that underpin individual motor behaviours and our complete motor repertoire. For instance, a hand grasp employs a different musculoskeletal network to wiggling an individual finger. Likewise, some motor behaviours engage more and less of our musculoskeletal system, for instance swimming, walking and cycling make more complex demands on our motor system than, say, blinking (see Gallagher, 2005, p.44-46.).

For the purposes of the ensuing discussion acknowledge Gallagher's claim that the body schema operates 'below the level of self-referential intentionality'. This means that the processes of the body schema are not directed towards the body in the same sense that an individual takes his or her body as an object of perception, say, when they wilfully move their body in some manner[32]. Instead, the body schema governs movements that are directed away from the body towards the environment. Think again of the example of your hand forming into a grasp in the right way for picking up the cup. In this case your motor behaviour was directed towards the cup. It is for this reason we are told in the passage above that the processes of the body schema are keyed in to the environment for the body schema would need to work out the dimensions of the mug to appropriate you hand into the right kind of grasp for picking up the cup. Objects of a different three-dimensional configuration would require a different formation of the hand.

The preceding considerations lead Gallagher to conclude that the body schema is best considered in terms of three functional sub-components. The first receives sensory information from the environment about the spatial configuration, position and movement of the body and objects in nearby space. Therefore, the first component processes sensory information that localises the position and relative spatial relations between the body and objects in its surroundings. The second component consists of what Gallagher calls *'intermodal capacities'* that 'facilitating the communication' between proprioception and the other sensory modalities, such as vision, about the position and movement of the body and surrounding objects. In other words, the second component somehow calibrates *multisensory* information about the disposition of the body and surrounding objects. The third component,

---

[32] This claim requires further qualification for surely the body schema can produce motor behaviours that are directed towards the body. For example, if I want to scratch my nose I may be aware of my hand approaching my face, but I may not always be aware of the activity of my finger as I bring my hand towards my face. Nonetheless, I think that what Gallagher is putting emphasis on here is the fact that the body schema is responsive to information from the environment and this is one key difference between the body image and the body schema (more on this point below).

the 'motor schema' facilitates bodily movement by recruiting all the informational resources that are required to activate the body in the way required to produce behaviour, such as the cortical and sub-cortical regions of the brain (Gallagher, 2005, p.46).

This much sums up the important aspects of Gallagher's line of thought with respect to the functional profile of the body schema insofar as it consists of multisensory processes that register and calibrate cross-modal information about the body and environment, and motor processes that are important for activating the musculoskeletal networks which are necessary for producing behaviours. It also lays down the foundations for his account of how the body schema can be differentiated from the body image insofar as its processes are prenoetic; involve the calibration of multimodal inputs from the body *and* the environment; and its processes are (mostly) directed towards the environment. To prove the reader with a full appreciation of these remarks I now turn to Gallagher's account of the body image, and his analysis of the relations between the body image and the body schema.

2.3 Gallagher's Functional Interpretation of the Body Image

Outwith academia the concept of the body image is the more familiar out of the two. To our common-sense understanding, the body image refers to the evaluative judgements and beliefs an individual may have about their body, such as the belief that their body is too short or too thin; too fat or too thin. We all have these kinds of beliefs and make such judgements about our body on their basis. There is a sense in which the scientific understanding of the body image concept is not so different from this folk conception. This aspect of the body image is known as the *body affect*. However, the scientific notion of the body image is more multi-faceted for it also includes the *body concept* which refers to the conceptual scheme that we use to label and identify our own body (i.e., 'my body'); its individual parts (e.g., 'arm'; 'leg' and 'head'), and ascribe to them a relative spatial location (e.g., 'left arm'). The most important aspect of the body image for our purposes is the *body percept,* which pertains to our personal-level percept of our own body. An individual's body percept may include information about the perceptible features of his or her body, such as the contours of his or her face; the relative relations between the different parts of their body, as well as the physical proportions of their body, etc. In other words, the body percept is best thought of in terms of a pictorial map of the body. It also includes motor generated sensory feedback. For

example, the qualitatively visual and proprioceptive experience an individual undergoes as she wiggles her fingers. The body percept also includes the experience of sensations on the body such as itches and pains[33].

In contrast to the processes of the body schema, the three distinct components of the body image make respective and collective contributions to an individuals' perceptual awareness, identification, and recognition of his or her body. For example, the aspect of my current body percept about a pain in my arm can only be an experience of a pain on *my arm* provided I am equipped with the ability to identify that the pained arm is part of my body, and not yours. Likewise, my current affectual feelings towards my body may be the product of the intricate interactions between my body percept and body concept. For instance, say I look at my hand which is identified as mine by my body concept, I may not like (my current affectual state) the way my skin looks (my current visual body percept) because of the inevitable consequences of the aging process. It is for reasons such as these that the body image is implicated in the development and transformation of the sense of bodily ownership over our own bodies; the sense of agency over our own actions, and our perception and understanding of the boundaries of our own body. In short, the processes of the body image play a key role in developing and shaping the different aspects of our perceptual apprehension of our body.

We are now in a position where we can draw up a provisional characterisation of Gallagher's interpretation of the distinction between the body image and the body schema. His account is consistent with the account from Head & Holmes on many fronts. For example, he defines the body schema as a sub-personal system of sensory processes that are, in part, implicated in registering proprioceptive inputs about the postural configuration and movement of the body. However, in the same vein as cognitive psychologists he concedes that the sensory processes of the body schema are multisensory and include information from the environment. The real distinctive mark of his account however is the additional motor component that actively produces motor outputs. When it comes to Gallagher's interpretation of the body image, he also includes the body concept and the body affect. Therefore, in accordance with the traditional characterisation of the distinction, Gallagher concedes that the processes of the body schema are sensorimotor and implicated in motor control, whereas the perceptual processes of the body image partly shape our perception of our body. Crucially,

---

[33] For more on what Gallagher has to say about the body image see Gallagher: 2005, pp.26-32).

the processes of the body schema and their behavioural outputs are not phenomenologically accessible, whereas the processes of the body image can be[34].

2.4. Empirical Evidence in Support of the Body Image and Schema Distinction

Having provided an initial characterisation of the body image and schema distinction based on their distinct functional components, Gallagher lends empirical credence to the claim that the body image and the body schema must be neurologically and functionally distinct by appealing to case studies which show that disruptions to the sensory systems which underpin one system don't affect the performance of the other, and vice versa. Recall from chapter one that this endeavour was already taken up by Head & Holmes as part of their attempt to lay out the distinction. In the same spirit, Gallagher offers two types of cases that demonstrate this same point. The first class of case studies pertain to deafferented patients who have a normal body percept but exhibit motor deficits[35]. The second class of case studies indicate that the opposite is also true; that is, an individual with a disrupted body percept may exhibit little to no motor weakness. One such case is case is unilateral hemi-spatial neglect (e.g., Denny-Brown, 1952) Because of a lesion or stroke an individual may only be able to consciously perceive one side of their body, but may still be capable of carrying out coordinated motor tasks. Their incomplete body percept is manifested in many of their behaviours, such as failing to dress one side of their body and only shaving half of their face, etc. Furthermore, this kind of lesion doesn't just engender an incomplete body percept; there is some indication that the effects of the lesion bleed into other aspects of their body image, as it is defined by Gallagher. For example, the individual may fail to recognise that a perceptually neglected body part belongs to his or her body. Feelings like the sense of ownership over one's body are an important aspect of the body affect. Nonetheless, the important point to note is that

---

[34] Gallagher rightly acknowledges that the processes of the body image need not be conscious all the time. Indeed, some of the underpinning processes of the body image may never arise to the personal level even though they may still directly shape our perceptual apprehension of our body in some non-trivial way. At different times, and in different contexts, our thoughts, beliefs and perceptions feature more or less explicitly our awareness of the state of our body, as well as the activities it is engaged in. The important point is that the processes of the body image, conscious or unconscious, play a direct role in shaping body perception. It is for this reason that the body image and body schema distinction does not reduce to the conscious/non-conscious divide, respectively. For more see Gallagher: 2005, pp.21-23.

[35] This kind of case has more far reaching consequences for Gallagher's analysis of the body image and body schema and we will discuss this type of case in the next section.

subjects who exhibit personal level neglect towards half of their body are still capable of using the neglected side of their body to produce motor behaviours that require left-to-right motor coordination, such as using both of their hands to tie up their shoe laces and walk in a coordinated fashion To a first approximation, this indicates that motor performance remains insulated from disruptions to the components of the body image which and this, in turn, lends empirical credence to the claim that the processes of the body schema are neurologically and functionally distinct from the processes of the body image[36].

By this stage of his analysis Gallagher has provided us with a clear conceptual dissociation between the body image and the body schema in terms of their respective functional properties (i.e., an account of what each system does), and the underlying functional components of each system. Furthermore, the cases of deafferentation and hemi-spatial neglect provide empirical support in favour of differentiating the processes of the body schema from the body image in this way. This conceptual framework offers a clearer idea of the processes that make up each system and the question is what we can then do with Gallagher's model. The available options, it would seem, are many. By prizing the body image and body schema systems apart in the way he does, Gallagher opens a window of opportunity to provide separate analyses of the body image and body schema. For example, one could probe more deeply into the details of the nature body image based on such questions as: what sorts of processes enable and shape our body percept, body concept and body affect? What defining and transformative effects, if any, do the causal interactions between each system have on our perceptual, conceptual and affective apprehension of our body? With respect to the body schema we may ask how do the 'intermodal capacities' calibrate cross-modal information about the body and environment; how the multisensory processes of the body schema work in conjunction with the motor processes to produce appropriate motor outputs? This list is by no means exhaustive, indeed the kinds of questions that we will be concerned with are questions about the nature of the relationship between the body image and the body schema, and how the relations between them are manifested at the behavioural level. The ensuing discussion shall be primarily interested with Gallagher's

---

[36] The case of hemi-spatial neglect, as I present it here, is not so clean-cut. Gallagher points out that some individuals exhibit a tendency to hop on the neglected leg some of the time and walk on it the rest of the time. Not too much hangs on this for the ensuing discussion. For a more comprehensive discussion see Gallagher, pp.31-34.

analysis of the relationship between the body percept and the body schema[37] in the following section.

2.5 Gallagher's Behavioural Analysis of the Interactions between the Body Image and Body Schema for Action through the Case of Deafferentation.

In the first section of this chapter I gestured towards the idea that the ingenuity of Gallagher's approach is that once he makes the conceptual separation between the body percept and the body schema, he then proceeds to work his way towards an account of the relationship between the two in the context of everyday action[38]. The guiding question of this part of Gallagher's analysis is why we need such a dual-system, each of which are equally capable of guiding and producing movement independently of the other. To address this question Gallagher appeals to a particular case of deafferentation which will be the subject of this section.

   The aforementioned considerations make it clear that the distinction between the body image and body schema is drawn up, in part, on the basis of their distinct functional roles. However, what Gallagher tells us is that the relationship between the two systems is much more intricate than that in the context of day-to-day, moment-by-moment, perceptual and motor activity. For instance, the postural positions and movements our body schema can influence our body percept. As a mature adult, we can walk without consciously thinking about how to put one foot in front of the other. Nonetheless, aspects of the motion of our body as we walk appear in our current body percept. Likewise, our current body percept can also play a role in guiding and monitoring our movements. For example, imagine that you are to thread, then you might visually guide the movement of your fingers to insert the thread through the eye of the needle. Clearly, we are not just experientially blind automata and so we know that we are not entirely dependent upon our body schema. If that's the case then why do we need a body schema to monitor and direct our motor behaviours below the level

---

[37] That's not to say that Gallagher overlooks or ignores the other aspects of the body image. In chapter 6 of his *2005* he acknowledges the need for more work to be done with respect to the importance of affectual processes for our mental and cognitive lives.
[38] We will see in the final chapter of this thesis that this is a very valuable aspect of Gallagher's analysis which hasn't been given the due recognition or attention it deserves despite its vital importance for developmental psychologists amidst their recent foray into the body image/body schema domain.

of consciousness, as well as a body percept that is implicated in the conscious guidance of movement? Why can't we just rely on a body percept? Why, then, do we need the two systems? Why do we not have just one or the other?

Gallagher's answer to the question of why we need a sub-personal body schema lies in his analysis of deafferentation (more on this below) and his answer to the question of why the body percept is implicated in the guidance of movement is provided in his behavioural analysis of the interrelations between the body schema and the body percept in the context of everyday behaviour. I will deal with each in turn. The first issue arises in the context of Gallagher's commentary in relation to the case of Ian Waterman (henceforth, IW), a deafferentated patient who lost all volitional control over his bodily movements because of flu-induced damage to his large myelinated nerves (e.g., (Cole:1995, Cole & Gallagher:1998). Ian lost the sense of touch, therefore he is unable to consciously feel, say, his hand gripping onto an object. He has also lost proprioception; the sense of the sensation of the position and movement of his body from below his neck. However, given that IW still receives visual feedback about the position and movement of his body and the interaction of his body in the environment; we can say that IW has retained some of his body percept that is anchored in what Gallagher calls *visual proprioceptive awareness*. Visual proprioceptive awareness is the visual feedback that's provided about the position and movement of his body which underpins an individual's qualitatively visual experience of the postural position and movement of the body and its parts. Visual proprioceptive awareness is distinct from the sub-personal visual feedback that's provided about the position and movement of the body. Gallagher calls the latter visual proprioception[39][40]. This notion of visual proprioceptive awareness is similar to what psychologist J.J. Gibson calls *visual kinaesthesis* which he defines in terms of the 'awareness of the locomotion of the body'[41]. Visual kinaesthesis can be active or passive. Active visual kinaesthesis refers to the active seeking out of information by the visual system about the movement of the body, such as when you track the movement of your hand that is holding the thread as it moves towards the eye of the needle. Passive visual kinaesthesis is the aspects of the movement of your body that you may be visually aware of, even though you do not actively seek such information, such as your visual awareness of various parts of your body as you walk. The key difference is that the former is

---

[39] This is an important distinction to keep in mind. I'll return to it later in this chapter and in chapter 5 of the thesis.
[40] For a more comprehensive discussion see Gallagher: 2005, ch.2, & Hurley: 2003.
[41] Gibson, 1968: p.341.

more likely to be implicated in the conscious guidance and monitoring of bodily movement than the latter.

Immediately after his illness IW lost all control over his motor skills. For instance, he was unable to keep his body and limbs in a fixed position and he was unable to instruct his body and limbs to move in determinate ways. On occasions when he would will his body to move in a certain way, the elicited movements would be unpredictable and uncontrollable. However, over an extended period IW regained most of his motor capabilities in the sense that he was eventually able to carry out the same tasks as he was capable of performing prior to his neuropathy. As such, he can lead a fully functional and independent life. However, to do so Ian needs to retain constant visual contact with his body and must constantly deliberate and monitor his movements at every stage of the movement trajectory. Therefore, to move his body effectively, IW is completely reliant on a constant supply of visual proprioceptive awareness. For instance, in darkness he is not able to control his movements and when he anticipates that he is about to sneeze he must sit down to avoid collapsing to the floor. IW must continually keep track (visually) of his body and his environment. To maintain his balance, he uses external objects as a fixed reference point and consciously tenses his muscles to freeze his body into its current position and to write he must constantly attend to his grip of the pen.

The important point is that visual proprioceptive awareness now undertakes the role that was previously taken on by the visual proprioception of his body schema (Gallagher, 2005, pp. 43-45). He later qualifies this statement by speculating that perhaps the motor programs which are usually available to the body schema to produce motor output remain unaffected and, through the rehabilitation process, have been re-accessed by the sensory components of the visual percept in the same way they are accessible on occasions when an individual recruits his or her body percept to consciously guide the movement of their body (Gallagher, 2005, p.47). Crucially, behaviours dependent on his body percept are not as effective as they were when they were produced by his body schema for his new motor routines are relatively slower and less precise, and they make more demands on his personal-level resources such as his attentional resources. Thus, IW reports that he is exhausted by the end of the day because of the concentrated effort that his day-to-day activities require.

The conclusions we are to draw from the case of Ian Waterman are threefold. First, it serves as another source of empirical evidence for the neurological and functional dissociation between the body image and the body schema. Second, having two functionally independent systems that can perform the same function in cases where the other is damaged

means that at least one of the systems is still available to allow an individual to perform the same task. The third is that having a fully functioning body schema which can produce and regulate our motor behaviours without the need for personal level intervention and monitoring *optimises* our behavioural performance in two important ways. First, the processes of the body schema make fewer demands on our personal-level mental and cognitive resources and the behaviours it produces are more effective and precise. The third revelation may seem surprising because it is very easy to underestimate the complexity of the neurophysiological underpinnings of even the simplest of motor behaviours. For instance, a single step recruits over two hundred muscles. Just imagine how long it would take us to walk if we had to consciously deliberate over how to retract and flex every muscle network from our toes to our head to keep our body and head upright whilst also deliberating about how to move each foot in front of the other. Thus, although I *know* how to walk in terms of the mechanics of the movement, I still do not know the exact muscle groups that are required to walk and maintain balance at the same time. Furthermore, it should come as a surprise to most of us that the orientation of the head tenses and relaxes the muscles surrounding the knee joints (Gallagher, 2005, p. 48). The take-home message is that having a body schema that can recruit the complex interconnected musculoskeletal networks required for our motor behaviours below the threshold of consciousness as adaptively advantageous because it relieves us of this burdensome and seemingly impossible task.

Nonetheless, there are two important exceptions to this general rule. Recall that the overhanging question of this section pertains to why we need a body schema and a body percept given that each system is capable of producing body movement independently of the other. The case of IW has illustrated the advantages of having an independent body schema. However, this doesn't respond to the issue of why we have a body percept. Gallagher addresses this issue in the context of his behavioural analysis of the relationship between the body schema and body percept during the process of motor learning, and when an individual must manoeuvre their body in hazardous conditions. I will explain the underlying rationale of each case in turn.

Gallagher's first claim is that the body schema and the body percept will be implicated in the day-to-day motor routines of a typical adult in complex, flexible and adaptive ways. However, during the process of developing their motor skills an adult is perhaps more dependent on their body percept than their body schema. The initial stages of the learning period are filled with awkwardness and uncertainty with respect to what one should do with their body. For example, as I am learning to shoot a basketball through a net I will pay closer

visual attention to, say, the force of my grip and the trajectory of my arm as it follows through to launch the ball towards the net. As I go I will learn to alter the force with which I throw the ball and readjust my aim in accordance with each miss. Meanwhile, other parts of my body will also readjust as I go along, such the force with which I push my knees down towards the floor before pushing up to get the right momentum for throwing the ball at the required speed. Therefore, motor learning is a process through which we continually develop and transform our motor repertoire and we do so through a trial-and-error period during which we consciously guide our movements and closely monitor the interactions of our body. Once we master these skills, so the story goes, they are then deferred to the body schema to be produced and regulated at the sub-personal level (see Gallagher, pp.45-54).

The same principle holds for the developing infant who is in the process of adding novel motor skills to her pre-existing skill set, albeit at a quicker rate. Through his analysis of neonatal imitation Gallagher tentatively suggests that some of the motor behaviours produced by the body schema are developed during gestation, such as opening and closing the mouth; sticking out the tongue, and bringing the hand to the mouth[42]. As such, a primordial body schema is already in formation when we are born. He then claims that, post-gestation, *learned* behaviours, such as sitting up; reaching and grasping; crawling, and walking are developed by a process through which the infant pays close attention to the movement of their body. For example, think of the little toddler who pays close visual attention to the wobbly movements of his or her legs as one foot moves in front of the other. After an extended trial-and-error period, which is fraught with unsteadiness and uncertainty, he or she becomes able to walk without needing to constantly pay close visual attention to their legs. In other words, the developing infant is, Gallagher speculates, more reliant on their body percept during the process of motor learning before those behaviours are deferred to the body schema when the behaviours can be carried out in a relatively stable and effective way, and this principle holds for motor learning in general under Gallagher's framework.

The next condition under which an individual may recruit their body percept in a more explicit way obtains when they must manoeuvre their body in potentially hazardous conditions, for instance when they are walking along an icy pathway or across a high ledge. In both situations, an individual must proceed with a degree of caution for one misstep or miscalculated move could cause a potentially painful or fatal injury. Therefore, in this kind of scenario they may use their attentional resources to plan and guide the movement of their

---

[42] See chapter 7 of his *2005.*

body by monitoring the progression of, say, their feet as they put one foot in front of the other. Thus, contra the case of IW, it pays to guide and monitor the movement of the body at the personal-level under these kinds of conditions.

The important upshot is that there are two conditions under which the body percept plays a more prominent role in the guidance and regulation of movement; namely, during motor learning and in hazardous conditions. Furthermore, Gallagher seems to be suggesting that this increased reliance on the body percept occurs when the likelihood and risk of making a motor error is increased and so it makes sense to elicit and readjust our bodily movements in accordance with personal-level feedback. In the case of motor learning this seems to be because the body schema is too underdeveloped with respect to the given motor skill. In the case of manoeuvring our body in dangerous conditions the implication(s) of making a motor error is increased because the outcomes of such mistakes carry a greater risk than normal. Thus, in situations where the rate of motor errors increase in an atypical way, and the stakes of such mistakes are higher, we are more reliant on our body percept for the guidance and regulation of movement.

2.6 An Analysis of Gallagher's Account of the Body Image and Schema Distinction.

At this stage let's just take stock of the key developments that have been outlined in the first half of this chapter. In response to the conceptual confusions that surround the distinction between the body image and the body schema, and to vindicate its status as a useful explanatory tool for cognitive scientists, Gallagher offers a conceptual dissociation between each system that is, in all essential respects, consistent with the traditional characterisation of the distinction. First, he defines the body schema in terms of a sub-personal system of multisensory processes which are, in part, involved in registering new inputs about the postural configuration of the body; second, he defines the body image as the system of processes that are involved in the perception of the body. Furthermore, his analysis of each concept also advances our understanding of the nature and functional role each system plays in the cognitive capacities they support. His analysis from the perspective of the phenomenology tells us that the body schema can produce motor behaviours which are out with the scope of our body percept. This reveals the second important point that because the body schema can produce behaviours independently of the body percept, it must consist of a

motor component that can directly produce bodily movements. Furthermore, the body image is much more than a sensory system that registers visual, tactile and motor induced feedback about the location of sensations on the body. This aspect of the body image, the body percept, also works alongside the body concept and the body affect to shape and transform our perceptual awareness and apprehension of our body. In general, Gallagher's analysis allows us to make a distinction between the body schema and the body image that is based on their different functional components. The body schema includes 'intermodal capacities' which facilitate the collaboration between the various sense modalities about the position and movement of the body and its relation to objects in peripersonal space, and a motor component which activates the parts of the musculoskeletal system required for movement and balance maintenance and is dependent upon gestational and post-gestational motor learning. In contrast, the body image is not only made up of sensory processes (i.e., the body percept), it is also made up conceptual processes (i.e., the body concept) and affective processes (i.e., the body affect) which underpin personal level perceptual recognition, identification and recognition of the body, as well as any evaluative feelings and beliefs which are directed towards the body.

Perhaps the most significant insights to emerge out of Gallagher's model of the distinction come from his behavioural analyses of deafferentation and the interaction between the body percept and the body schema in the context of motor learning and everyday action. With respect to the case of deafferentation, the important conclusion is that having a body schema that can elicit and regulate motor behaviour independently of the body percept optimises our motor performances because its operations make fewer demands on our cognitive resources, and its outputs are more effective and precise. With respect to motor learning and everyday action, the first important conclusion is that the relationship between the body percept and the body schema is much more delicate than the coarse conceptual dissociation between each system seem to suggest. The body percept has an important role to play in shaping and facilitating the motor repertoire of the body schema during motor learning, and is also recruited on occasions where an individual must manoeuvre his or her body in potentially hazardous conditions. Ultimately, an individual employs his or her body percept when the risk and rate of error is increased; either because his or her body schema is underdeveloped or because of the current situation that they find themselves in. Thus, under these conditions the body percept plays an equally important role in the optimisation of motor behaviour.

In sum, Gallagher offers a clear framework for understanding how the distinction should be carved up and a clearly defined set of issues about the importance of the body image and

body schema for the production, transformation and optimisation of our motor skills. Nevertheless, it also seems to me that his analysis throws up as many questions as it answers[43]. I will outline the explanatory gaps that I intend to fill in the remainder of this thesis. The first set of questions pertain to the nature of the body schema. For example, how is cross-modal information calibrated (i.e., what is the 'intermodal capacity'? and how does it function?), and how does this information figure in the development, transformation and selection of motor output? The second set of questions pertains to the relationship between the body percept and the body schema and its significance for the development, transformation and selection of our motor output. First we might ask what is the nature of the processes through which the interactions between the body percept and the environment shape the motor component of the body schema? Second, if the likelihood of error explains the alternation between the body percept and the body schema, what role, if any, do those errors play in determining and transforming our motor repertoire? Another way of putting this same question is how does the minimisation of error figure in the optimisation of the behaviours that are produced by the body percept and the body schema? [44]. This list of questions is by no means exhaustive, but these are the questions that I will content myself with in the remaining chapters of the thesis. Ultimately, a comprehensive model of the role that the body percept and the body schema play in the production and optimisation of motor behaviour must provide a rationale for addressing such questions.

2.7 The Argument against Gallagher's Account of the Body Schema from De Vignemont

Gallagher's analysis of the body image and schema distinction is perhaps the most widely acknowledged. However, to the best of my knowledge there is only one existing critical response to his treatment of the body image and schema distinction in the philosophical literature and it comes from Frederique de Vignemont as part of her short review of his *2005*[45]. Therein Vignemont analyses several important strands of thought that feature in his book. Here I will concern myself with her critique of Gallagher's definition of the body

---

[43] This remark is not intended as a criticism of his framework for I will argue that the response to his framework in the philosophical literature has only served to stifle the plethora of questions that his analysis opens up.

[44] I attempt to address the former set of questions in chapter 5 of the thesis and the latter in chapter 6.

[45] De Vignemont, 2006.

schema. Recall the passage from Gallagher that's quoted in section 2.2 above wherein he defines the body schema as a sub-personal system of sensorimotor functions, capabilities, etc., none of which have 'the status of a conscious representation'. Vignemont finds both aspects of this claim problematic for quite distinct reasons. I will rehearse and discuss each in turn. With respect to Gallagher's conceptualisation of the body schema in terms of sensorimotor functions and the like she says:

> 'One may first notice that Gallagher defines the body schema as a set of capabilities. I have some difficulty understanding what these capabilities are and Gallagher does not provide any help on this topic. He remains very vague about the components of the body schema' (De Vignemont: 2006, p.3).

I think there is a way of reading this line of thought which is correct. I mentioned earlier that perhaps it would be more effective to describe the body schema in terms of the set of functions that underlie motor behaviour. This is quite consistent with Gallagher's general functional description of the components of the body schema, and it is on this point that I disagree with the latter point in the passage from Vignemont. She argues that Gallagher remains vague with respect to the components which make up the body schema. But in section 2.2 of this chapter I explained that Gallagher clearly lays out what he takes the functional components of the body schema to be. The first component registers cross-modal inputs; the second component (i.e., the so-called 'intermodal capacities') calibrate the information that's provided by the first component and the third component is the motor component that produces motor outputs by activating the musculoskeletal underpinnings of our behaviour patterns. Thus, it is not quite accurate to say that Gallagher is vague in his description of what he thinks the body schema consists of.

Nevertheless, her critique of his statement that the body schema does not have the status of a conscious representation is perhaps the more important of her objections to his account. Gallagher tells us in the passage quoted above that the body schema is not a conscious representation. This statement seems uncontroversial because the body schema is inaccessible to personal-level consciousness. De Vignemont's argument is not that the body schema must be a conscious representation; rather, she claims that Gallagher's non-representational analysis of the body schema is explanatorily impotent in virtue of the very fact that he omits

all talk of internal representation from his conception of the body schema. This is peculiar, she says, for 'it seems clear that we do need information about the state of our own body' to move our body and this, she claims, 'argues in favour of pragmatic body representations for action' which are implemented in the brain and are part of the body schema (De Vignemont, 2006, p.3).

The natural question at this stage is what Vignemont means by a 'pragmatic body representation for action'. To explain, Vignemont draws upon a distinction between higher-order visual representation and pragmatic visual representation of physical objects (Jacob & Jeannerod, 2007). A higher-order visual representation of an object represents what you can do with the object, in other words the range of actions that can be taken up with respect to it. For instance, a knife can be used to saw, slice and chop various objects. Thus, the functional property of an object puts constraints on what an individual can typically do with it[46]. In contrast, a pragmatic visual representation represents the perceptible features of the object that are relevant for action, such as its size and its shape. This puts further enabling conditions and constraints on the ways in which the object can and cannot be used. For example, its shape will determine how it can be grasped and its size may determine whether an individual will need to use one or both of their hands to pick it up. Together, higher-order and pragmatic representations play an important role in determining the scope of possible actions that can be performed with respect to an object based on its functional properties and its objective configuration.

At the next step of her argument, Vignemont transposes this distinction to the body schema by conceptualising it in terms of a higher-order and pragmatic body representation. The higher-order body schema represents a 'functional map' of the motor repertoire of the 'human body in general'. This amounts to the idea that the higher-order body schema represents the range of actions that the human body can perform in virtue of its anatomical configuration. It also represents the constraints of the human body for 'what is true of my kinaesthetic constraints is normally true of your kinaesthetic constraints' (De Vignemont: 2006, p.6). This representation also represents the more or less economical and efficient ways of carrying out certain tasks. In contrast, the pragmatic (or 'first-order') body schema represents the perceptible features of the body and its parts; such as the spatial configuration of the body, and the size and shape of the body parts. Importantly, the latter also consists of occurrent information from the sensory systems about the postural position and movement of

---

[46] Clearly this is not always the case. I can, for instance, use a knife as a sort of make-shift screwdriver and use it to injure someone. Nonetheless, the general principle still holds.

the body is processed by a forward model (which she also ascribes to this component of the body schema) which uses occurrent inputs for comparison with recent inputs to predict the future postural positions of the body based on its current movement trajectory.

By construing the body schema in this way, Vignemont sees is that her model has more explanatory leverage than Gallagher's. However, I think this argument rests on two unwarranted assumptions. The first assumption is that by refusing to conceptualise the body schema in terms of a representation Gallagher overlooks the obvious fact that, as she says, we need information about the current state of our body to move. This is clearly not the case as he ascribes this function to the first and second component of his tripartite account of the body schema. The other side of the coin is her assumption that by construing the body schema in terms of representation her model possesses explanatory leverage that she thinks Gallagher's interpretation lacks. I find this assumption more problematic for it is not at all obvious that conceiving of the body schema in terms of internal representation imbues her conception of the body schema with some sort of magical explanatory power that she seems to think non-representational accounts lack. For instance, it's not at all clear to me that 'first-order body representation' has a kind of explanatory power that 'intermodal capacity' lacks insofar as both are *described* as the integrative processes by which multimodal information is received and calibrated. I italicise the word 'described' here intentionally for I think that Vignemont is confusing a description for an explanation, and if Gallagher's account is problematic, in her view, because it is descriptive and not explanatory, then it seems to me that her account faces the same problems, at least until she probes more deeply into the nature of the representations in question that would allow us to understand the functional properties she ascribes to each representation[47]. My overarching point is that a functional description of the body schema doesn't quite add up to a comprehensive explanation of the neural basis of multisensory processes in the brain until it is supplemented an additional story about how they carry out their ascribed function. Nonetheless, it helps get such a story off the ground.

---

[47] I provide a hypothetical response to this objection in chapter 5 of the thesis.

2.8 De Vignemont's Conception of the Body Schema in Terms of Forward Modelling, Bayesian Sensorimotor Prediction and Action-Oriented Representation.

Insofar as her critique of Gallagher's conceptualisation of the body schema goes, I think Vignemont's model of the body schema is problematic for reasons I have already expressed in the preceding section. Nonetheless, this isn't the only interpretation of the body schema that Vignemont offers for in her *2010* she proceeds to develop her account of the body schema further. Importantly, this analysis goes some way towards filling the explanatory gap I just mentioned with respect to the idea that a functional description of the body schema that's construed in representational terms doesn't add up to an explanatorily superior model. In a similar vein to Holmes & Spence, Vignemont puts the general Bayesian model of sensorimotor processes forward as a plausible theoretical framework for understanding the nature of the body schema. Furthermore, she re-conceives of the body schema in two distinct ways. The first offers a reading of the body schema in terms of a predictive inverse-forward model. The second offers a reading of the body schema in terms of action-oriented representation. The first key argument of this section is that the first reading raises important questions pertaining to the importance of Bayesian style sensorimotor prediction for motor control, whereas the second reading is problematic on the grounds that her definition of action-oriented representation is too restrictive and her analysis of the body schema, so defined, regularly confuses the body schema for the body percept.

First, let's consider her conception of the body schema in terms of a forward model. As we will see in more detail in chapter 4 in the thesis the concept of a forward model is an important explanatory concept for outlining the computational structure of the informational processes that underpin the planning, production and re-modification of motor output. The following account from Vignemont should provide the reader with a clear enough idea of what a forward-model entails. Vignemont provides an interpretation of a forward model in terms of three distinct aspects of the body schema. First there is the *initial body schema* which computes the motor commands necessary to move the body to a desired state, relative to its current position and its movement capabilities and/or constraints. The second component is the *predictive body schema* that 'anticipates the sensorimotor consequences of the body' and 'allows for the anticipatory control of movement'. Then there is the *updated body schema* which 'carries sensory information only about the bodily parameters that have

changed' on the basis of the instructed movement (De Vignemont, 2010, p.672)[48]. Thus, the operations of the initial body schema occur prior to motor output and presumably the predictive body schema is in operation prior to (and presumably during) the motor activity. The updated body schema is in operation once the movement trajectory unfolds and provides a continual flow of sensory information that's fed back to the predictive schema.

The interesting part of this analysis is that it raises the question of the importance of sensorimotor prediction for the planning and execution of motor control. This is an aspect of the body schema that Gallagher's analysis does not consider. Recall that Head & Holmes interpret the body schema in terms of a 'combined standard against which all subsequent changes of posture are measured before they enter consciousness' (Head & Holmes, 1911-12, p.187). This is another way of saying that the body schema is a forward model of sorts. However, the real distinctiveness of Vignemont's conception of the body schema is her emphasis on the importance of the interactions between top-down and bottom-up processes for determining the outputs of the body schema. To get a handle on these two important issues she appeals to the Bayesian model of sensory processing. Recall from chapter 1 that the relevant top-down processes are probabilistic predictions about the sensory consequences of bodily movement which are brought to bear on motor induced sensory data.

However, Vignemont doesn't quite get as far as addressing the question of how Bayesian style sensorimotor predictions shape and facilitate the operations of the *body schema* for what she provides is an explanation of how the *body percept* is enabled by Bayesian predictions. This story is established in three steps. First, she claims that whatever the body schema is, it must be conscious. A compressed version of the story goes like this. The body schema may not require consciousness to be able to perform its function, but that doesn't necessarily mean that the outputs of the body schema do not enter consciousness. She appeals to motor imagery as an example. For instance, If I close my eyes and form a visual picture of my body moving in my mind without actually moving my body and it is the body schema that is responsible for movement related processes, then this visual imagery of my body in movement must be produced by my body schema. This can't be caused by my body image because that is only responsible for producing a percept of what my body looks like, not its movement. Thus, the

---

[48] The crucial difference between this account of the body schema and the account of the body schema Vignemont provides in her *2006 is* that the 'initial body schema' is equivalent to the 'higher-order representation'. Furthermore, the 'predictive body schema' and the 'updated body schema' are equivalent to the first-order body schema'. The Important difference is that her updated account conceives of the three distinct aspects of the body schema as belonging to one cohesive forward model which is involved in simultaneously eliciting motor commands, predicting their sensorimotor consequences and providing relevant sensory feedback.

story goes, 'the body schema must be conscious' and thus 'the availability to consciousness is *not* a criterion to differentiate the body schema from the body image' (De Vignemont, 2010, p. 673). However, Vignemont proceeds to use such criteria to claim that the body-schema is a particular species of representation: action-oriented representation. By her definition:

> A body representation is action-oriented if and only if it carries information about the bodily effector (and the bodily goal in reflexive actions) that is used to guide bodily movements (De Vignemont, 2010, p. 672).

This means that the processes of the body schema are exclusively directed towards the body. Thus, the body is both the performer of actions as well as the intended subject of those actions. The line of justification is this:

> 'The body as a goal does not seem to differ from any other bodily goals. I can reach for
> for my head or I can reach for my book, and it seems *strange* to think that the body
> schema represents the book. When I reach for my book, my movement is guided by
> a visuo-motor representation of the book that recruits the dorsal system. Similarly,
> it makes sense to assume that when I reach for my head to scratch it, my movement
> is guided by a proprioceptive-motor representation that recruits the body schema.
> (De Vignemont, 2010, p. 672).

The plausibility of the claim that it is 'strange to assume that the body schema represents the book' rests on what it means for the body schema to represent something. The succeeding statement would suggest that the notion of representation Vignemont is working with consists of a contentful percept of the book. Whether this is true is an open question. However, considering Gallagher's analysis we have sufficient reason for thinking that the body schema can, and indeed does, direct motor responses that exhibit a sensitivity and responsiveness to the configuration of our environment. Think again of the example of reaching out for your coffee-cup whilst your hand forms into a grasp outwith the scope of your awareness. This would suggest that the conception of the body-schema in terms of action-oriented representation that is directed only towards the body is too restrictive. It also goes against the

grain of the standard view of action-oriented representation[49]. This brings me to the next point when Vignemont claims 'my movement is guided by a visuo-motor representation that recruits the dorsal stream'. The flip side of the coin of Gallagher's analysis is that conscious visual guidance is only responsible for the guidance of the more coarse-grained aspects of the reaching movement. The body schema takes care of the rest (i.e., the more fine-grained motor adjustments and the maintenance of balance). As such, there is no reason to think that *only* visual guidance is at play here. The key question is why, in this scenario, the visual guidance of the movement of the arm is *not* the outcome of the processes of the body schema, but yet motor imagery *is* the product of the body schema? In other words, why is my imagining seeing my body move the outcome of my body schema but my seeing my body move not? What's the difference between an imaginary and a bona-fide visual percept of the moving body?

The source of the problem for Vignemont is rooted in the failure to respect the boundaries between the body percept and the body schema, i.e., whether their outputs are directly available to consciousness, which she initially rejects based on the case of motor imagery, but then uses to justify the claim that the processes of the body schema are exclusively directed towards the body. The case of motor imagery doesn't rule out the possibility that it is the body percept that is responsible for this mental exercise. Perhaps it could even be said that motor imagery is an offline reconstruction and/or recombination of the relevant aspects of the body percept. For example, when I close my eyes and visualise, in my mind, my body moving in some way, my mental vision has similar phenomenological content to my veridical visual experiences. Likewise, a physically impaired individual could, with the right amount of mental effort, form a similar mental image.

2.9. How to Make Progress on the Philosophical Analyses of the body schema and its relation to the body Percept: A Concluding Summary

The significance of this line of criticism I have levelled against Vignemont's characterisation of the body schema in terms of a conscious action-oriented representation is that it only serves to reinforce the importance of Gallagher's model if philosophers are to avoid falling

---

[49] Millikan, 1983, Mandik, 2005. More will be said about how action-oriented representations are characterised under predictive processing framework in chapter 4.

into the same pitfalls as cognitive psychologists. The implication for her model is that her account of the body schema in terms of a forward-model that subscribes to Bayesian driven sensorimotor prediction fosters a productive avenue of research that the contemporary analyses of the body schema do not consider, yet this has ultimately led her to a dead end because we are still no closer to a Bayesian interpretation of the body schema and Vignemont now concedes that the distinction ought to be dropped on the same familiar grounds as Holmes & Spence who say that the distinction just causes too much confusion and no consensus can be found over how these concepts should be dissociated and analysed.

In response, we can say that representational analyses in cognitive psychology and philosophy are equally responsible for that for note that they always take the same course and end up with the same outcome. A definition of the body schema is provided in broad representational terms, however the conception of the body schema in such terms is too liberal to respect the intended boundaries of the distinction. As such, the confusion persists and the diagnosis is that the problem must lie with the distinction. A different outcome thus requires a different approach. The question is what such an approach should look like if we are to make any headway. The approach I will take up will preserve (most of) the unproblematic insights associated with the representational analyses of the body schema whilst respecting Shaun Gallagher's non-representational analyses of the body schema and the body percept/body schema distinction. In what follows I will pursue the currently abandoned line of thought from Vignemont's analysis of the body schema in terms of an action-oriented representation that is guided by Bayesian sensorimotor prediction and, by doing so, I will take the necessary steps towards providing a reconciliatory approach to defining the body schema which neatly straddles the pro- and anti-representational divide.

Chapter 3

A Bayesian Perspective on the Optimising Principles of Cross-Modal Cue Integration

Introduction

The analysis of the issues pertaining to the body schema and body schema distinction across chapters one and two make it clear that multimodal sensorimotor processes optimise our behavioural interactions in the environment. Furthermore, Shaun Gallagher's analysis of deafferentation and the relationship between the body image and the body schema in the context of everyday action showed that this is true of the sub-personal processes of the body schema which work independently of the body percept. This is also true of the body percept during motor development and hazardous conditions. To begin to build our story about how to conceptualise the body schema and its relationship with the body percept, we first need a more fundamental story about the optimisation strategies which maximise the effectiveness of the outputs of perceptual and motor processes. In this chapter, I will outline and examine empirical evidence from Marc Ernst and his interpretation of the data from the perspective of Bayesian causal inference which suggests that sensory information isn't just passively received and recapitulated into given outputs.; rather, the brain evaluates and integrates cross-modal signals based on their task-specific and context dependent reliability. The story to emerge is that this process optimises the output signal in two ways. First, calibrating the signals improves the accuracy of the output; second, flexibly weighing up and recalibrating the signals in accordance with their task- and context-relative strengths and weaknesses under different conditions ensures that the outgoing processes remains stable and unperturbed by disruptions to the signals. This will provide the foundations required for the ensuing line of thought that will be developed in the subsequent chapters of the thesis.

3.1 Bayesian Models of Perception: The Inherently Unstable Nature of Sensory Processing and Why Context Matters

In chapter one of the thesis we got a brief flavour of the central tenets of the Bayesian vision

of sensory processing (e.g., Bernardo & Smith (1994) &. Gelman, 2004). In this chapter I

will further explore how Bayesian causal inference has been used to establish a general

theory of perception. To better understand the motivations of the Bayesian picture let's begin

by considering the implications of two problems the brain must overcome to ensure the

perceptual outputs it produces are accurate and stable. In his *2006*[50] Marc Ernst considers the

---

[50] Ernst: 2006, pp.1-5.

question of how the perceptual outputs of sensory processing remain stable in the face of inherent disruptions to the feeding signals that are brought about by the inevitable noise that's caused by sensory signal transfer. Sensory processing, by its very nature, inevitably causes random fluctuations to the information that's carried by the incoming signals. As such, the information provided by the unending flow of new and updated signals can never fully accurate. As such:

> […] no information-processing system, neither technological nor biological, is powerful enough to 'perceive and act' accurately under all conditions.
> (Ernst & Bulthoff, 2004, p.162).

Furthermore, the severity of the degradation a signal undergoes is also proportionate to the conditions under which the sensory signalling is taking place for information provided by the sensory modalities may be stable and accurate in some conditions, and not in others. For example, as I type this chapter in broad day light I can see the features of my visual surroundings in clear detail. This is because vision is a very reliable source of sensory feedback in normal lighting conditions. However, in complete darkness I may not be able to see very much at all. I may be able to identify the contours of some objects around me, but the visual information I am getting is not as reliable as the information I am provided with in broad daylight. It certainly isn't reliable enough for me to completely rely on it to get myself out of a dark room safely. Instead, the best thing to do is to feel my way around the room to feel for objects to avoid a collision and perhaps feel for a wall or a door to feel my way out. In this case, touch is more reliable than vision. However, there are some occasions in which vision is more reliable than touch. For example, we have all had the experience of our fingers going numb from holding onto a freezing cold object, such as an ice lolly, for so long that we can no longer feel our fingers.

Against this backdrop, the guiding question is how the brain manages to synthesise the sensory information available to it from its perceptible surroundings into a coherent, accurate, unified and world-reflecting percept to compensate for the degradations to the signals. What's interesting is that from a first-person perspective we are not aware of the fluctuations to the sensory signals, however the fact that the brain does compensate for the context-dependent fluctuations to the signals is manifested in our behaviour. In the scenario just

outlined, we increase our reliance on the available signal and decrease our reliance on the unavailable signal. The fundamental issue that must be addressed is how the fluctuations are registered and how they are compensated for at the neurological level in a way that helps us make sense of the changes in our reliance to the signals.

3.2. The Maximum Likelihood Principle

To address this question Ernst proposes what he calls the *maximum likelihood principle* (or MLEP) which claims that the brain utilises prior-beliefs ('priors') to evaluate the incoming sensory data and calibrates *redundant* cross-modal signals based on the current quality of the contributing signal to produce the most accurate estimation of the causes of the input that it can produce (Ernst, 2006, p.6). Under MLEP perceptual outputs are, at best, estimations of the cause of the stimulus. A clear illustration of this comes from Green & Swets[51] who showed that if an individual is asked to provide one hundred estimations of the size of an object, all one hundred estimates will be different. Thus, perception is a *probabilistic* process that provides estimations of what our perceptible surroundings are like[52].

3.3 Empirical Evidence for the Maximum Likelihood Principle

To empirically verify the claim that cross-modal signals are processed in accordance with MLEP Ernst appeals to earlier studies of his (e.g., Martin Banks & Marc O Ernst (2002)) to explore the possibility that the perceptual system does maximise the accuracy of perceptual outputs by combining cross-modal signals to provide a combined estimate that reflects the current quality of the signals. The general methodology the authors apply is one we are already familiar with considering the rubber hand illusion experiments which were discussed in chapter one. The methodological strategy seeks to create competition between a visual and a haptic (i.e., a sensor input that is provided by touch) signal about the comparative size of

---

[51] Green & Swets: 1988.
[52] See Ernst: 2006, pp.2-5 for a more comprehensive overview.

two objects by deliberately increasing noise to one signal to see if the weighing up of the signals flexibly recalibrates in accordance with the MLE principle. If this is the case, the weight of the tactile signal should increase proportionally to the decrease in availability of the visual signal.

During the task the participants had to estimate the comparative size of paired objects. The perceived objects were paired in terms of the time delay between them, with shorter time delays between the presentation of each object that were part of a pair and longer time delays between each pair. As the paired objects were presented in sequential order the participant had to estimate whether the objects in each pair were the same or different in size (see Ernst, 2006, pp11-13). For test purposes, there were three conditions under which the objects could be perceived. Under the first condition the objects could only be seen; under the second condition they could only be touched, and in the third they could be simultaneously seen and touched. Using Bayesian inspired sophisticated mathematical statistics[53] Ernst & Banks worked out the respective initial estimation of the visual and haptic signals and the respective estimates each provides about the comparative size of each set of paired objects. For test purposes, they worked out how the individual signals were weighted in the seen only and felt only conditions in response to paired objects with no difference in size. What's interesting is that in the simultaneously seen and felt condition, the final estimation was a combination of the visuo-tactile signals which was weighted slightly in favour of the visual input. This suggests that the default response of the brain is to combine the signals (more on this below).

Using the principle that interference to the dominant visual signal should systematically reduce its influence to the weighting of the combined signal and increase the weighting of the less dominant signal, Ernst & Banks systematically introduced noise to the visual signal (Ernst, 2006, pp. 8-11). The upshot is that in accordance with their mathematical calculations, the weight of the visual signal did systematically decrease and the influence of the haptic signal increased proportionate to the decrease in the weight of the tactile signal. When the input from the visual signal was reduced to 0% it was completely overridden by the tactile signal[54]. Under the MLEP this would be expected because additional noise reduces the salience of the information from the interfered signal. In turn, this increasingly undermines the accuracy of the information it is providing. With decreasing accuracy comes a decline in the reliability of the signal and this means that its assigned weight gets smaller and smaller. This then tips the scales, as it were, in favour of the other available signal(s) whose weight

---

[53] The mathematics isn't necessary for my purposes here. See, e.g., Bram:2013, ch.3)

[54] Recall from chapter one that this phenomenon is called sensory capture.

gets proportionally higher to compensate for the interruptions to the affected signal. And when one signal is unavailable, the other takes over (Ernst, 2006, p.11).

The importance of these studies is threefold. First, the default response of the brain in response to signals that report on the same property or event is to bring the signals together to provide a combined estimate; second, the weighing up of the signal is a function of a) the availability of the signals and b) the current qualify of the inputs they provide; third, when one contributing signal is completely unavailable the other signal takes over. The real beauty of the studies from Banks and Ernst is that for the first time the processes through which the brain evaluates and recalibrates sensory signals to maximise the accuracy of its output despite the inherently instable and fluctuating nature of sensory processing can be mathematically quantified and examined under test conditions.

3.4 The Implications of the MLE Principle: Understanding the Optimal Significance of Redundant Cross-Modal Signals

The natural question we can ask is what insights we can glean from the empirical studies and the MLEP with respect to our interest in the optimal significance of cross-modal signalling. The first thing we can say is that the default response of the brain is to always combine the sensory signals to compensate for the inherently noisy nature of sensory signalling. Interestingly, the weighting of the initial combined signal weighed slightly in favour of vision. On reflection, this is unsurprising for visual input is the most reliable modality for reporting the spatial dimensions of three-dimensional objects. For example, I can judge the size of an object just by looking at it. However, if I close my eyes and feel the object I can't judge its size just by holding it. This doesn't mean that I can't come to an accurate judgement about the size of the object by touching it; rather, vision provides the same data much more quickly than touch. Thus, the first way in which the brain optimises the multimodal inputs that underpin perceptual outputs is by combining the signals in accordance with their context-dependent reliability and efficiency. In turn, this entails that the brain must possess prior knowledge about the context-sensitive reliability of the modalities.

The second important point pertains to the optimal importance of having multiple available sources of information about our perceptible surroundings in the first place. At the initial

stages of this chapter, the fact that the brain has multiple sensory channels from which it can glean information about the world was cast up in terms of a problem in the sense that the more sensory information is processed by the brain makes sensory processing noisier. However, what the studies show is that having multiple sources of complementary sensory information is adaptively advantageous for the brain can exploit their respective, context-dependent strengths and weaknesses to maximise the accuracy of the driving signal. For instance, if I was in the dark room and only had a visual system I would have no other reliable resources to rely on to manoeuvre my body safely around the room and locate the exit. The MLEP principle neatly explains what is occurring in this situation. In normal lighting conditions the weight of our visual signals about our surroundings has a higher weight than tactile information. However, in sub-optimal lighting conditions the weight of the visual signal decreases and the tactile signal increases because the optimal thing to do is increase your reliance on the system that is available and providing a steady stream of sensory feedback.

What is under discussion here is the optimal significance of *redundant* sensory signals (Ernst, 2006, p.5). Sensory signals are redundant when they report on the same property or activity in the world, i.e., they carry equivalent information in their respective formats. Non-redundant signals carry different information, although they can carry distinct information about the same property or event (e.g., a visual signal about the colour of the mug and a tactile signal about its shape.) The notion of redundancy in the multisensory integration literature is quite distinct from the common-sense understanding of this term. In ordinary discourse, the word redundant usually means that something which is redundant is needless or is no longer of use or value. The connotations of the word redundant as it is used in the context of multisensory integration are important for understanding the significance of the way redundant cross-modal signals optimise the mental, cognitive and behavioural capacities they support.

The working notion of redundancy comes from systems engineering. In systems engineering the components of a system are redundant when there are two or more duplicates of a particular component part of a system. Typically, whilst one component is active the redundant components sit in reserve, but are readily poised to override the functioning component if it undergoes a malfunction to some degree. Without duplicates this could lead to the system performing less than optimally or failing to work at all if one of its essential components has failed. For instance, to prevent a complete system failure most modern aircraft have redundant systems built in, including duplicate flight controls, electrical

generators and dual-engines. Having redundant systems in place which can perform in the exact same way simultaneously prevents against operational failure and ensures that the system works to full capacity without major disruption(s). Another way to put the point is that having duplicate systems which can perform the same functional role in the system is adaptively advantageous for it one fails, the other can replace is so that the operations of the system can carry on as normal, or at least as close to normal as possible.

If we transpose this line of thought to the sensory systems, the same argument applies to the optimal significance of redundant signals, only in this case it is not only one sensory module that does the driving work whilst the others are inactive waiting to be activated in instances where the driving signal fails. Rather, each sensory signal is active and makes a continual contribution to the outgoing signal. Of course, the extent to which the signal is contributing to the output is dependent upon the current quality of the information it carries. Moreover, as the studies from Banks & Ernst suggest, if one signal is so degraded to the extent that it is providing no input, the others) can override and the perceptual process can go on as normal. This is the way the brain can maintain stability in the face of the inherently unstable nature of its own sensory processes and the context-dependent reliability of the sense modalities. As such, Ernst expresses the importance of redundant signals thus:

> 'Redundant signals' may to some degree sound like a waste of information. But actually this is not necessarily so. There are two major advantages in having redundant information available: the first is that the system is more robust, because when one estimate is not available at a given time (or its information is degraded) the other estimate can substitute for it. The second advantage is that the final estimate becomes potentially more reliable compared with the reliability of the individual estimates feeding into the combined estimate. (Ernst, 2006, p.7)

Importantly, combining the signals to maximise the accuracy of the signal increases the efficiency of the process for the individual by decreasing the amount of variance (i.e., any discrepancy) between the perceptual estimate and the actual property being reported. In other words, the other side of the coin to maximising accuracy is minimising error (Ernst, 2006, p.9). This is an important point because the prior likelihood that the integration between signals is likely to produce a high volume of errors determines the integration relationship between them (Ernst, 2006, p.11). When the relations between cross-modal signals are likely to produce a higher risk of error they are combined which is to say that they are temporarily

coupled together for the duration of the estimation process. This entails that the brain still has access to the combined signals to recalibrate the signals in accordance with the varying quality of the signals. In contrast, when the signals present a lower risk of generating error they are 'fused' which is to say that when the signals are combined the brain no longer has access to the individual signals (Ernst, 2006, p.19).

In support of this claim Banks & Ernst[55] studied the integration relationships between redundant and non-redundant signals. Using a similar experimental set up they compared the coupling relations between non-redundant signals within the same modality (visual-only signals) and redundant cross-modal signals (visual-haptic signals). What they found was that the integration between non-redundant visual signals about binocular disparity and texture was fused, whereas redundant visual and haptic signals about size were combined (Ernst, 2006, p. 19).

In general, it makes sense for the brain to retain access to both signals when they are only coupled for a temporary duration, such as during tool use. Up until the point at which I pick up my rake to collate freshly cut grass the maximum reaching distance of my hand is constrained by the canonical boundaries of my body. However, once I pick up the rake the reaching distance of my hand instantaneously extends proportional to the length of the rake. As such, motor planning needs to be updated to quickly re-adapt my reaching behaviour accordingly. In contrast to this, when the coupling between signals does not, and need not, change quickly and their integration presents a low risk of error, then access to the individual signals need not be retained. In other words, if it is sufficient for related errors to be resolved over the long term, then the brain doesn't (always) need access to the independent signals (Ernst, 2006, p.21). Clearly, the experiments from Ernst suggest that the integration between visual signals about disparity and texture is one such case. In low risk cases where the consequences of error are not serious, repeated exposure to the same property over time is all that's required to gradually update the estimate (Ernst, 2006, p.20).

---

[55] Banks & Ernst: 2002, Ernst, 2006, pp.19-24.

3.5 Reinterpreting the Rubber-Hand Illusion in terms of the MLEP

Ernst's MLEP provides a rationale for understanding how the relations between cross-modal relations optimise perceptual outputs. In general, cross-modal signals are weighed up in accordance with their context-dependent reliability and are continually calibrated and recalibrated during the perceptual process in response to fluctuating incoming signals. In other words, the more reliable the information provided by a sense modality is under current conditions, and throughout the process, the more the brain increases the perceptual systems' reliance on it to improve the efficient and accuracy of the output; ensure that perceptual processing remains stable, and minimise error. This puts us in a nice position to explain some of the empirical data that we considered in chapter one with respect to the ability to accurately track the location of sensations on the body. One such example was the rubber hand illusion. In this task, the participants see and feel tactile stimulation on their own hand. However, the crucial difference is that the hand they see receiving tactile stimulation is a rubber hand while the hand they feel tactile stimulation on has been occluded. When asked to locate the stimulus, the participants say that they feel the stimulus on the rubber hand. The guiding question is why vision trumps proprioception and tactile inputs about the location of the hand and the stimulus, respectively. From the perspective of MLEP, the visual signal is not coupled with the tactile signal for the visual signal is reporting about the visual location of the stimulus in space (which is directed towards the rubber hand). In contrast, the tactile and proprioceptive signals are redundant for they are reporting on sensory feedback from the actual hand. However, vision is the modality on which we usually rely on to locate objects in space and so the prior likelihood that vision will deliver a more accurate estimate than touch causes an increased reliance on the visual signal. In this case, the visual signal completely overrides the tactile/proprioceptive signal[56].

---

[56] It is not quite clear at this stage why we see complete sensory capture in the rubber-hand illusion from this perspective. However, I will try to explain this from the perspective of the prediction-error minimising paradigm.

The significance of Ernst's MLE Principle: A Concluding Summary

The overarching purpose of this chapter is to provide an illustration of one of the most important works in the multisensory integration literature about the way in which multimodal processes optimise the perceptual processes they underpin. What Ernst's empirical data provides us with is an empirically grounded principle (i.e., MLEP) for thinking that the default optimisation strategy of the brain is to combine redundant signals to maximise the accuracy of the driving signal that determines perceptual output. In turn, the availability of multiple sources of sensory information that can fulfil the same functional role in the perceptual system means that perceptual processing can carry on as normal if one of the signals becomes degraded and/or unavailable. In this case, one sensory channel can override the other (malfunctioning) modality. This is just part-and-parcel of the brains' ability to flexibly recalibrate the relations between the signals in accordance with their context dependent reliability. In turn, the more accurate the signal the less likely it is to be incorrect. As such, the MLEP provides an account of one of the processes through which the brain compensates for the instability of sensory signalling and invariantly reliable signal sources to enhance its perceptual outputs. However, this can only be part of our story. For one thing, a story still needs to be told about *how* the relevant prior knowledge that's brought to bear during the evaluation and calibration of sensory signals is built up. Furthermore, to get to grips with a Bayesian model that will do the job of providing a framework to interpret the body image and body schema, we still need a story about the optimal significance of multisensory integration for sensorimotor integration. This will be the focus of the next chapter.

Chapter 4

Predictive Processing: Uniting Perceptual and Motor Processes under Active Inference and Free

Energy Minimisation

Introduction

There is a relatively newer and updated version of the Bayesian paradigm that was outlined in chapter three which has been gaining pace in cognitive science: the prediction error minimisation model (PEM). In this chapter, I will look at the central tenets of this framework as they apply to perception and action. Under this model, the optimal existence of biological agents requires an internal, hierarchical model that can successfully predict the internal state changes (i.e., sensory stimulation) that are brought about the external causes of the stimuli (in the case of perception) and its actions. So understood, perception and action are united in an effort to contribute towards the optimal existence of the agent as a whole by continually testing out the predictions and making relevant revisions (i.e., minimising and/or eliminating 'prediction errors') in order to enhance the predictive accuracy of the model. What we will discover is that this model provides all the necessary resources we need, and more, for understanding how multimodal processes optimise perceptual and motor processes, and offers a unifying story about how perceptual and motor processes work together to optimise the interactions of agents in their environment.

4.1 The Inferential and Predictive Nature of Perception: From Helmholtz to Enactivism.

The prediction error minimisation paradigm (PEM)[57] takes its lead from the Bayesian model of perceptual inference that was considered in the previous chapter in several key ways. This model is, first and foremost, a theory about the structure of information processing in the brain for it adopts the principal Bayesian idea that bottom-up sensory inputs are met with top-down prior beliefs. However, it offers a distinctive view of what kind of mechanism the brain is on a whole. In a similar vein to Ernst's Bayesian perspective, this framework starts with the premise that the brain is an optimisation device that seeks to maintain stability and maximise the accuracy of its outputs. However, the strategy by which it does this is not by combining incoming signals to produce a combined estimate; rather, the task of the brain is to accurately

---

[57] E.g., Hohwy:2007, 2012, Friston *et al*: 2008, 2009, 2010, 2011, Clark: 2012, 2015, 2016.

predict the causes of the stimulus before the stimulus arrives at the periphery, and it does so by bringing descending multi-level predictions from cortical and sub-cortical regions to bear on the incoming sensory data to try to work out the cause of the stimulus[58].

As a model of perception PEM adopts the general idea that was outlined in the previous chapter about the nature of perception in terms of knowledge driven inference. However, its starting point provides a more comprehensive understanding of why we should think of perception in this way. To draw out the central tenets of the PEM model, I will take this aspect of the theory as my starting point. The standard way of thinking about the nature of the mediating transformative processes that transform sensory data into perceptual outputs, such as the view espoused by Marr (1982) for instance, take perceptual processes to be the processes by which incomplete, inverse, two-dimensional retinal inputs are gradually transformed into a complete, world-reflecting, three-dimensional image that is perceive in experience. On this view, each level of the perceptual hierarchy makes its own specialised contribution to the transformation of the perceptual signal at a particular level of detail and scale in isolation from the others, before passing it along to the level above. The process continues until the process is complete and the sensory signal is transformed into the final percept.

However, there is a consensus these days that perception involves more than the transformation of available sensory inputs because sensory data does not suffice to provide the rich and seemingly detailed experiential content of personal level perception. One reason for this concerns the consideration we encountered in the previous chapter, namely that there is no one-to-one correspondence between an input and its cause for an infinite amount of causes can generate the same input, e.g., the input pertaining to colour, like red, could be caused by a red ball, a red car, a red rose, a red tomato etc. and the one output can bring about a multitude of effects[59].

Another consideration comes from Alva Noë[60] in relation to what he calls the 'problem of perceptual presence' where aspects of your surroundings may be 'present' in your experience but absent from the sensory feedback that the sensory modalities are currently providing. Cases of amodal perception are a good example of this phenomenon. To explain, let's begin with the epistemologist's favourite example of looking at a tomato that's sitting on a windowsill. Suppose that you are looking at the tomato. It seems you as if you are

---

[58] For more see Clark: 2013, pp. 181-190
[59] See Hohwy: 2012, ch.1.
[60] Noë: 2004.

experiencing, visually, a whole, three-dimensional, fleshy, juicy fruit but all you can see is the incomplete surface properties of the tomato in virtue of its position relative to your field of vision. The same applies when you look at a chair tucked under the table for it doesn't seem to you as if your experience of the chair is in any way incomplete, yet all you are really seeing is the top of the chair. Hidden from direct view is the back of the chair, its legs and the seat. Thus, what we perceive at the level of experience outstrips the sensory data that's currently available (see Noë, 2004, pp.67-69). This means that there must be additional mediating processes that provide the additional information about your sensory surroundings that you experience and this indicates is that the processes which occur between the uptake of sensory data and perceptual output that are involved in determining the content of personal level experience involves more than just the impression and transformation of data on the sensory apparatus. Instead, it involves making inferences about what's 'out there' in the world on the basis of the current sensory data you are receiving. Alas, we have the first tenet of PEM which is inspired by the theory of perception from Helmholtz (1878) according to whom perception involves unconscious inference. Broadly speaking, he claims that sensations merely represent the stimuli that cause them; they are not exact copies of them and the learned correspondence between sensations and the perceptible objects that produce them is built up through interactive experience.

Questions about the nature of the sort of knowledge that underpins the mediating inferential processes of perception have been central to sensorimotor and enactive approaches to perception (e.g., Noë & O'Regan, 2001, Noë, 2005), according to which the content of experience is given by the possession and exercise of a kind of sensorimotor knowledge which is a practical knowledge of the systematically causal relationship between the movement of the body and the sensory apparatus. For example, in the case of looking at the tomato, the amodal content of my experience (i.e., the content which I cannot directly see) is given by my understanding that if I were to move my body in precise ways relative to the tomato, say by changing my position relative to it, other bits of the tomato would come into view. It's this understanding which gives me the impression that I'm looking at a whole, three-dimensional tomato and not a façade of a tomato (e.g., see Noë, ch.2, 2005). What's distinctive about the enactive view is that perception is inherently anticipatory and that the sensorimotor knowledge which underpins a perceivers' grasp of the effects of movement on sensory stimulation is not just causally important to perception, in the same way that the availability of light is for instance, *it is constitutive of what perception is*. Without it, perceivers wouldn't have the kind of perceptual capabilities they have (Noë, 2005, p. 2).

The PEM model adopts the Helmholtzian line of thought that perception is, by its very nature, inferential and is underpinned by an understanding of the causal correspondence between objects and the sensory stimuli they generate and the enactivists claim that the knowledge which underpins perception is inherently anticipatory.

4.2. Hierarchical Perceptual Processing

Under PEM the perceptual knowledge which underpins our perceptual access to the world around us is part of a probabilistic 'generative model' in the brain of perceiving agents, so called because it generates predictions about the correspondence between sensory stimuli and the objects which cause them. The generative model is hierarchically structured with higher levels eliciting more abstract perceptual and lower-levels which elicit predictions about more fine-grained aspects of perceptual detail. At each level, there is a 'representation unit' which elicits the predictions which pertain to the probable causes of sensory input at the specific level of perceptual detail that that level pertains to and an 'error unit' which receives error signals from the lower levels in the hierarchy (e.g., see Clark, 2012, p.21-32).

What is distinctive about the PEM model, Clark tells us, is the story it has to tell about the communication that occurs between levels in the hierarchy. On the traditional view, the signal itself is passed up the hierarchy. Under PEM, *only the error signal* is relayed up the hierarchy until it is cancelled out. This occurs when one level in the hierarchy can account for the erroneous detail. If this fails to happen a change in generative model may be in order. Thus, what gets fed downwards through the hierarchy are the predictions. On the standard view, each level works in isolation and makes its own contribution to the transformation of the signal before passing it up. Under PEM the representation unit at one level can feed into the level below to try to cancel out a prediction error via 'forward connections'. Furthermore, the only signals which are relayed in an upwards fashion from low to high levels in the hierarchy are the error signals. In the case of perception this process allows for the rapid processing of perceptual data, and the multi-level interactions within the perceptual hierarchy allow for the quick detection and correction of prediction errors as any conflicts between the predictions elicited under a given model and the inputs can be resolved at an earlier stage of perceptual processing (Clark, 2012, p.54).

This means that when generative models are employed to explain incoming sensory data, the information that is being worked over is the information that is predicted given the current model that is under use. Different generative models will thus carry different representational units that carry different predictions and replicas of the perceptual data. Thus, it is argued that it is not raw sensory input that contributes to the final percept but is instead a 'virtual version' (Clark, 2013, p.13) of that data.

What this gives us is a clean explanation of how top-down and bottom-up processes work in tandem to cancel out prediction errors. The selected generative models provide a measure against which the detection of errors is made possible for prediction errors can only be generated when sensory evidence is evaluated under a given model, and the alternation between models helps cancel out a wide array of prediction errors. Within each model the multi-level communication between top-down representation units play an important role in cancelling out prediction errors at multiple scales of perceptual detail. This enhances the capacity for a given model to cancel out prediction error more successfully for each level in the hierarchy will be equipped to resolve discrepancies at different levels of description. It also enhances the capacity for prediction error minimisation at each level if each representation unit doesn't work in isolation. This means that in principle the cooperation between each level can cancel out any discrepancies on the spot instead of being passed up the hierarchy for further processing.

It is for this reason that HPP engenders a 'gist-at-a-glance' approach to perception to explain the importance of top-down influence (Clark, 2013, p. 14). The idea in mind is that when the brain receives a given sequence of sensory data it employs a generative model that best captures the more global attributes of the initial inputs. This may include the more global aspects of the distal scene, such as the dense expanse of different bodies in a vast crowd. However, as more specific sensory feedback comes in more specific inputs may be registered, such as the colour of the clothes that people are wearing, their gender, whether they are fighting or dancing, and so forth. The adjustments that are made as the incoming sensory sequences unfold will depend on how much information is accurately captured by the current model, i.e., how much of the visual scene can be accounted for by the representation units. If any discrepancies are left unresolved another model might be recruited whose predictions are better equipped to account for the incoming data, or perhaps the current model will be re-altered so that it is better able to account for the data the next time that it is deployed. This means that the importance of top-down and bottom-up interaction within the hierarchy is revealed in two different ways. Firstly, sensory input is required to get the

prediction error minimisation process started and unfolding sensory feedback provides the basis for providing prediction errors. In turn, initial sensory data is met by a selected generative model to explain its cause and further amendments to the top-down predictions that are generated are driven by the prediction errors which are generated as the sensory trajectory unfolds.

4.3 Perception as Prediction: In summary

In summary, the PEM framework characterises perception as a process by which a generative model is issued in response to incoming sensory inputs to cancel out prediction errors. This is initially facilitated by the multi-level communication between each level in the hierarchy of the selected generative model. In cases where the prediction error cannot be cancelled out by the initially selected model, a direct switch may be made to an alternative model or the alternation back and forth between two or more models may occur to cancel out the error and, in doing so, perception optimises the predictive capacities of the brain by making appropriate internal changes in response to its inputs. Therefore, the final percept is determined by the interactions, and settled agreement, between each level in the hierarchy of the generative model that has successfully proven to account for most of the details presented in the current visual scene. Thus, PEM differs from the orthodox model of perception in two fundamental ways. The first concerns what the perceptual system is doing for under this framework the job of the perceptual system is not to reconstruct the sensory data into a percept, but to cancel out prediction errors. Secondly, it suggests that the communication between levels is multilateral, in other words the predictions at one level can influence the performance of the level below to cancel out the error signal.

4.4. Action-Oriented Predictive Processing: Casting Motor Control in terms of Prediction Error Minimisation.

A similar story has been mounted by Karl Friston & his colleagues about the hierarchical nature of the motor system and the computational underpinnings of motor processing to explain the nature and functional role of motor control (e.g., Friston *et al.* 2008, 2009, & Friston, 2010, 2011). Andy Clark recently labelled this framework the 'action-oriented predictive processing' (henceforth AOPP) model of sensorimotor functioning[61]. In a similar vein to the story about the nature of the perceptual system, the motor system itself is characterised as a hierarchical system that implements *sensorimotor* predictions about the sensory consequences of movement to cancel out perceptual and motor generated prediction errors. The authors take as their first step the fact that the motor system is hierarchical in nature. For instance, the regions of the brain that are involved in more abstract motor planning, i.e., the pre-motor and motor cortex, are in cortical areas like the precentral and postcentral gyrus, whereas the regions involved in the ongoing retuning of motor output in response to incoming signals, like the cerebellum, are linked closer to the spine and muscle networks to produce quick motor readjustments[62]. Thus, the more stable parts of our action routine are situated further away from the spine and those involved in quick motor responses are closer to the effectors.

The distinctive hallmark of AOPP pertains to how motor signals are characterised in terms of sensorimotor predictions. This is quite a radical claim for it states that having sensorimotor predictions about the sensory consequences of action is both necessary and sufficient for bringing action about:

> We only need to have expectations about the sensory consequences of moving
> in order to elicit movement (Friston *et al* 2010, p.1.)

Friston *et al* (2010, p.3) use the example of how the fulfilment of expectations about tactile experience drive an agent to act. To explain the line of argument the authors use an example

---

[61] Clark:2013.

we have encountered already where an individual is trying to feel their way around a dark room to turn the light switch back on. Feeling ones' way around the room is directed by the expectation that objects can be located by touch and on this basis an individual may probe their way around to find the light switch try to confirm those expectations, and will continue to do so until they are proven to be correct (Friston, 2010, p. 3). In this case body movement in the room is directed by, and aimed towards, the fulfilment of certain expectations about the tactile feedback that should be experienced as an individual feels their way around the room. In the case of action, prediction errors drive a change in motor output until the desired expectation (as dictated by the generative model) has been satisfied. Thus, prediction errors engender a change in motor output to bring about sensory data that fits with the generative model, not a change to the selected model to fit the data.

Mobile agents are capable of actively generating and regulating the streams of incoming sensory stimuli that their sensory apparatus is met with on the basis of their implicit understanding of how their body movements bring about changes to sensory inputs in desirable ways, for instance a shift in eye gaze brings previously unseen parts of the world into view and straining one's eyes brings fuzzy objects into sharper focus. In this case, simple motor adjustments at one level (i.e., the sensorimotor level) bring about the anticipated changes to the sensory stimuli that satisfy an individual's expectations about what he or she will see at another, i.e., at the level of personal level experience. To explain the thought in mind let's suppose that a motor command is recruited to instruct the hand to reach outwards on a given trajectory until it meets a target object, like a cup. However, say that the retrieving hand fails to come into contact with the cup once it has reached its maximal reaching distance. This will generate a prediction error that will be corrected for by employing another motor strategy, like leaning the trunk of the body over to some degree to extend the reaching distance of the hand so that it can come into contact with the cup. In this example, prediction errors instruct the motor readjustments the body makes until the prediction elicited by the generative model has been satisfied.

Of course, this is only part of a much bigger computational process that recruits the perceptual and the motor system. Recall that whatever the motor system does, it does so in conjunction with information about the location of objects relative to the body. Proprioceptive information simply isn't enough to confirm whether an action has been successful, for instance confirmation that the hand has successfully grasped the mug might be confirmed by visual, tactile and proprioceptive information. As such, AOPP posits that the perceptual and motor system work together under an operating principle called *active*

77

inference where the perceptual system elicits predictions about the state of the body and the world and engages the motor system to actively sample the sensory evidence in accordance with the brain's generative model(s). Thus, we are told:

> In this view, the central nervous system is not divided into motor and sensory systems but is one perceptual inferential machine that provides predictions of optimal action, in terms of expected consequences. (Friston *et al*, 2009, p.9)

This point needs some clarification as it is potentially misleading. There are no strict lines of division between the sensory and motor system under Friston's version of AOPP because the nature of the motor signals are themselves sensory in nature. Nonetheless, that's not to say that there is no distinct perceptual and motor system. On the contrary, their divisions are carved up based on the predictions they specialise in. The perceptual system elicits predictions about the causal relationships between states of the body and environment and the sensory stimuli they produce. The motor system elicits sensorimotor predictions about the sensory data that's generated by the motor behaviours of the agent. Nevertheless, this is not to say that they work separately; rather they are unified under active inference to cancel out prediction errors[63].

Under active inference, both perception and action are 'in some deep sense computational siblings' (Clark, 2013, p. 19) because they subscribe to the same computational profile of employing hierarchical generative models to cancel out sensory prediction error even though the way they do this is quite different. This is reminiscent of the line of thought about the relationship between perception and action from psychologist J.J. Gibson:

> Each perceptual system orients itself in appropriate ways for the pickup
> of environmental information, and depends on the general orienting system
> of the whole body. Head movements, ear movements, hand movements, nose
> and mouth movements and eye movements are part and parcel of the perceptual
> systems they serve. They serve to explore the information available in sound,
> mechanical contact, chemical contact and light. (Gibson, 1966, p.58)

---

[63] See Clark, 2012, 2015 for more.

The overarching point is that PEM and AOPP accords a different job to perception and action respectively. The job of the perceptual system is not to reconstruct an internal model of the external world, but to retain correspondence between the predictions elicited by the generative model about the causes of sensory stimuli and the external objects they represent. Likewise, the job of the motor system is not just to produce motor outputs but to do so in a way that confirms or disconfirm anticipated sensory feedback which is specified by the generative model those expectations. If prediction errors are produced, then it is the job of the motor system to readjust the outputs until the expected (sensory) outcome has been achieved. As such:

> Sensory prediction error controls movement and top-down predictions control sensory prediction error (Friston, 2010, p.2)

In summary, under AOPP perceptual and motor processes prescribe to the same computational process of bringing to bear hierarchical generative predictions to bear on incoming sensory data to cancel out sensory prediction errors. The perceptual system does this by adapting the model to fit the data and the motor system does this by bringing about necessary changes to the inputs that will better fit the expectations that have been specified by internal generative model(s). As such, perception and action work together to suppress prediction errors under *active inference*. So far this story doesn't go quite so far as to explain how it supports the optimal interactions of the agent. This will be the topic of the ensuing sections in this chapter.

4.5. Reconceiving the 'Forward Model' in terms of the Hierarchical, Prediction Error Minimising Scheme.

The central argument in defence of AOPP is that it provides an improved understanding of the functional structure of a forward model. This concept features briefly in chapter two regarding Vignemont's interpretation of the computational underpinnings of the body schema. On her account, an inverse model (i.e., her '*initial body schema*') computes motor

commands on the basis of the intended movement and priors about movement constraints a forward model (i.e., her *predictive body schema*) which elicits sensorimotor predictions about the sensory consequences of the selected movement, and the *updated body schema* provides sensory feedback for comparison with the predictions elicited by the predictive body schema. In contrast, Friston's conception of the forward model is much simpler and thus offers a view of the motor system which is more computationally tractable, or so he claims. Under his interpretation, forward models only consist of sensorimotor predictions and prediction errors; there is no need for a dual inverse model for the sensorimotor predictions are already in the business of predicting the sensory consequences of movement and prediction errors take on the job traditionally ascribed to efference copies (i.e., their presence or absence provides the sensory feedback necessary to affirm or disconfirm that the motor prediction has been fulfilled)[64]. Optimal behavioural routines are thus selected for based on their ability to fulfil the sensorimotor predictions elicited by the motor system and are built up by active inference, not on the basis of their expected utility or reward in relation to the agent's current adaptive goals and intentions[65].

At this stage, we are in a better position to understand precisely why Andy Clark dubs this model of motor control in terms of prediction error minimisation the 'action-oriented' framework. The notion of action-oriented representation is not new (e.g., Mandik, 2005). Ruth Millikan's 'pushyu-pullmi' representations are a good example of action-oriented representation. Pushyu-pullmi representations consist of so-called 'descriptive content' that specifies, or describes, the current state of the world that is important for the guidance for action, such as the spatial position of objects relative to the body and 'directive content' which is a motor command (or sets of motor commands) that instruct the body to move based on the information encoded by the descriptive content. The dual content of pushyu-pullmi representations is both perceptual and motoric and is very much at the service of directing motor behaviour in the environment, hence why such representations are said to be action-oriented (Millikan, 2008)[66].

If one accepts that motor processes involve a kind of action-oriented representation then one can tell an important, and unified, story about the nature of their content and how that content is updated and maintained in terms of active inference. The descriptive content of action-oriented representations is encoded in a generative model which elicits probabilistic

---

[64] See Friston, K *et al* 2009, 2010 and Clark, A & Pickering, M, 2014 for more.

[65] This line of thought will become clearer in the section below on the 'free-energy principle'. Challenges for this claim will be raised in the final section of this chapter.

[66] More on this in chapter 6.

perceptual predictions about the state of the body and world, and the directive content is constituted by sensorimotor predictions that elicit movements on the basis of the descriptive predictions. Furthermore, the descriptive and directive content of action-oriented representations work in conjunction under active inference to cancel out prediction errors to optimise the interactions of the behaviours of the agent within its environment.

According to Clark there is another added advantage of explicating motor control in terms of the kind of action-oriented representations that are posited by the predictive processing framework. In a nutshell, he claims that such representations avoid the negative stigma that's usually associated with 'mental representation' *tout court* which usually imply that an internal model of the world and the body is internally reconstructed every time there is a change to the state of the body and the environment. In contrast, the representational content of action-oriented representation is already encoded at various scales and levels of detail at each level in the hierarchal and the only information that is directly worked over, and might have a role to play in reconfiguring the content of the model, is provided by the prediction errors. The point is that the bulk of the work conducted by the representational contents of action-oriented representations takes place before inputs impinge on the sensory apparatus. This is very nicely explained by Andy Clark thus:

> The brain, in ecologically normal circumstances, is not just suddenly 'turned on' and some random or unexpected input delivered for processing. So there is plenty of room for top-down influence to occur even before stimulus is presented. This is especially important in the crucial range of cases where we, by our own actions, help to bring the new stimulus about. In the event that we already know we are in a forest (perhaps we have been hiking for hours) there has still been prior settling into a higher representational state (Clark, 2013, p. 23)

Clark's point is that more often than not, the activities we engage in do not require sudden and drastic changes from our internal economy that underpin sophisticated perceptual and motoric re-adaptations; exceptions to this rule might include walking out of an enclosed building into wide open space or taking part in a triathlon which requires a drastic change in motor performance in water and on land within a short period of time. Thus, the bulk of motor processing is driven by relevant sensory information in the form of sensory predictions about the state of the environment and sensorimotor predictions about the movement of the body that *act as a proxy* for the actual information that should be provided by the body and

environment in normal circumstances. Thus, the bulk of the work involved in structuring and maintaining action-oriented representations is top-down heavy.

4.6. The Free Energy Principle

The broader issue that emerged out of the earlier chapters of this thesis concerns how we can explain how the body image and schema, and the relationship between the two, supports the optimal motoric interactions of an agent in his or her environment. What this first requires is a story about how perceptual and motor processes optimise behavioural interactions in the environment. This is important for at this stage the reader may question why casting up perceptual and motor processes in terms of action-oriented representation is relevant to this endeavour. Indeed, why frame perception and sensorimotor processing in terms of prediction error minimisation in the first place? The answer to this is given by a particular part of the AOPP framework under consideration which supports the idea that there is an important link between prediction error minimisation and optimal behaviour, and it is explained in terms of how prediction error minimisation at all levels of neural functioning is at least a necessary condition for optimising the chances of survival of biological organisms as a whole.

This line of thought comes from Friston in the context of a particular reading of AOPP in terms of the so-called free energy principle (e.g., Friston, 2007). AOPP need not always be understood under this principle (e.g., Howhy, 2012), however it provides a clear explanation of why the adaptive success of biological agents is essentially dependent on the ability of their nervous system to make accurate predictions about the consequences of incoming sensory input. The central claim of the free-energy minimisation principle is that for an individual agent to optimally thrive in its environment, each state change in its body must minimise free energy. Roughly speaking free energy is the upper bound on the (predicted and variable) state transitions that an agent will undergo as a result of the interactions of its embodied states. This includes everything from cellular and neuronal activity all the way to the interactions of the body in the physical and social environment (Friston, 2007, p.1). In other words, the principle states that if the world is in state X, then the agent is likely to be in state Y. Likewise, it states that if the agent performs some action X, it will cause internal state change Y, given the current model and anticipated sensory evidence. The free energy of a biological agent is a function of its sensory inputs and what Friston calls a 'recognition' or

'ensemble' density which is a model of the formal constraints on the motion of the agent's states (i.e., as each state transitions from its current state to the next) given its morphological requirements. As such, the recognition density is:

> […] specific to each class of agent. The ensemble density can be regarded as the probability of finding an agent in a particular state, when observed on multiple occasions or equivalently, the density of a large ensemble of agents at equilibrium with their environment. (Friston, Dainizeau *et al*, 2010, p.3)

In effect, the recognition density is characterised as the provider of the best probable guess on the most likely cause of a given internal state of the agent given the current internal state of the agent and incoming sensory data. Likewise, the recognition density also encodes the effects that the agents' alterations to, and actions upon, its environment bring about in terms of internal state change. Free energy is the difference between expectation about cause of input as specified by the recognition density and the actual input. Therefore, the more the sensory input complies with the expectations of the recognition density the more it minimises their free energy. On their formal specification, the more that each and every state change strives to minimise free energy the more they increase the long-term entropy of the agent, i.e., its state changes, perceptions and behaviours. In turn, this maximises its resistance against disorder.

Thus, the underlying rationale of the free-energy principle is that the survival of an animal depends on its having an internal model of how its internal states are affected by the states and activities of its environment:

> Not only does the agent embody the environment, but the environment embodies the agent. This is true in the sense that physical states of the agent (its internal milieu) are part of the environment. *In other words, the statistical model entailed by each agent includes a model of itself as part of that environment. This model rests upon prior expectations about how environmental states unfold over time. Crucially, for an agent to exist, its model must include the prior expectation that its form and internal (embodied) states are contained within some invariant set.*
> (Friston, 2011, p.89, my emphasis added)

In a nutshell, the free-energy principle puts constraints on the behaviours of the animal for the higher the likelihood that an agent (or its brain) can accurately predict the changes that it will undergo as a result of changes to its embodied states which are produced by its activities at all levels of description before they occur, the less energy it has to consume readapting to unexpected changes. In turn, this ensures a higher probability that it can withstand (i.e., survive) those changes to its embodied states.

Thus, the adaptive success of different classes of biological agents and their individual members is directly dependent on their ability to satisfy their basic homeostatic requirements, e.g., secure food, shelter, warmth and procreate. In their pursuit of these basic requirements they must be able to regulate their internal states within certain parameters to sustain their physiological integrity. This means that changes to its internal states must be kept within morphological bounds. In short, for an individual agent to exist in a more than sub-optimal state it requires that its embodied states – i.e., activities and changes at the genetic, molecular, cellular, neuronal, and bodily level do not compromise its physiological integrity.   This provides the first line of justification for the first claim that an agent must have a model of itself as part of its environment. The second line of justification is explained via the role of the recognition density for generating expectations about how environmental events will unfold and bring about internal state changes to the agent to ensure that the activities of the agent are constrained within certain bounds. Being able to maintain a state of equilibrium in its internal milieu during ongoing state variation presupposes that an agent (or at least its brain) can predict the consequences of those changes.

4.7. The Importance of the Free Energy Principle: A Brief Summary

In sum, the importance of prediction error minimisation for optimising the activities of biological agents is best understood from the perspective of the free-energy principle. It states that the optimal existence biological agents requires that their internal state changes minimises their free-energy which, in turn, ensures that their internal activities respects their morphological constraints. One way of ensuring that internal state changes stay within appropriate bounds is to be able to predict the consequences that are brought about as the internal economy of the agent transitions from its current state to the next. So far, this has

been explained under the free energy principle in terms of an internal model (the recognition density) and current sensory input. In response to the current input the recognition density elicits predictions about the most likely cause of the input (under the current model). Free energy is the difference between the *predicted* and the *actual* cause of the sensory input. Thus, the more the state changes of the agent minimise free energy the more precise the predictions about the consequences of its behaviours become. This characterises the adaptive significance of the importance of sensory related predictions under the free-energy minimisation principle. By striving to minimise free energy at all levels of description, from cellular interactions to motoric behaviours, the brain attempts to minimise its own predictions errors about the causes of its sensory inputs.

4.8. Understanding the Optimal Significance of Multimodal Integration in terms of Free Energy Minimisation and Active Inference.

If the biological brain really does constrain the activities of the agent on the basis of the free-energy principles then it gives us a better grip on the role that prediction error minimisation plays in the selection and refinement of the interactions of the agent at the behavioural level because, if this is the case, behaviours are selected for and built up on the basis of their capacity to support the optimal interactions of the agent in its environment. This means that not just any old behavioural strategy that will fulfil the predictions elicited by the generative model will be selected; rather, optimising behaviours are those which have predictable consequences under the agent's recognition density to minimise potentially hazardous and/or fatal consequences.

Importantly, Friston *et al,* 2010 attempt to demonstrate this line of argument to show how multimodal processes underpin active inference support optimal motor behaviours. Using a sophisticated system of Bayesian mathematical algorithms they constructed simulations of different generative models that predict the state of the environment and instruct the behaviours of the 'system' that are governed by free-energy minimisation. In one experiment, Friston and his colleagues attempted to simulate a generative model that mimics the visuo-motor system and a visual object. The purpose of this model was to instruct the simulated

visuo-motor system to track the movement of the visual object in the same way that the eye traces an object so that it stimulates the central foveal region of the retina. The guiding question is whether priming the model with different prior expectations that would provide the basis for creating prediction errors about the movement of the stimulus would influence the saccadic behaviour of the eye[67].

At the first stage the model was primed to expect that the object would remain stationary. In turn, the stimulus was programmed to remain in the same position so that the expected and actual movement of the object were the same. In other words, no prediction error should be generated in under this initial condition and as expected no movement of the visuo-motor system was observed. At the second stage an intentional prediction error was generated by priming the visuo-motor system to expect that the object wouldn't move and programming the object to move. Under this condition, the prediction error was cancelled out by the activity of visuo-motor system which 'tracked' the movement of the stimulus. In this case, the generated prediction error compelled the movement of the eye to cancel it out.

The explanation the authors provide for this result is that when the foveal region of the retina is fixed on an object it provides consistent and stable stimulation to the retina. Moving objects come in and out of the foveal region and as such provide sporadic and unstable retinal stimulation. One way of stabilising the perception of the object in the visual field so that it doesn't appear to pop in and out of visual awareness is to visually track the object so that it continually stimulates the centre of the retina. Shifting the gaze in accordance with the movement of the object cancels out the prediction error because it provides a continual and steady supply of visual data to the retina in the same way as stationary objects do when they eye is fixed on them. The simulation provides an example which is analogous to the way in which retinal stabilisation can cancel out movement induced prediction errors. This might not be a true reflection of how retinal stabilisation is achieved in all ecologically realistic conditions, but not enough is built into this specific simulation to account for this fact. For example, background motion may play an important role in facilitating retinal stabilisation but the specific experimental set up here is silent on this point for the activity of the simulated oculomotor system is directed towards a visual target. Here is the explanation Friston *et al* provide:

---

[67] See Dazineau, Friston *et al*, 2010, pp.12-19.

> Orienting or tracking behaviour […] enables the perception of a moving
> target as if it was stationary. Although the target may be moving, the full
> extent of this movement is not registered by perception. This is because the
> prediction error…is explained away by action, and action is not part of the
> perceptual model. In other words, from the moving point of the agent, the
> target is not moving. (Friston, Daunizeau *et al*, 2010, p. 13).

In other words, the visuo-motor activity of the eye elicited by the oculomotor system cancels out a prediction error that has been generated by the perceptual system which it could not cancel out on its own. At the next stage of the experiment, the generative model was primed to expect that the object would move towards a target area. In this case the object remained in the same position such that a prediction error is generated when there is a shift in gaze but the object does not move in tandem and this is precisely what occurred. The gaze shifted in the direction of expected movement before returning to back to the visual target. In this case the initial eye movement is guided by the expectation about movement of the stimulus, and cancelled out by moving the gaze back to the stimulus (Friston *et al*, 2010, pp. 14).

The key upshot of the simulations, the authors argue, is that they provide a very simple demonstration of how priors about the sensory consequences of movement produce initial behaviours and how prediction errors induce corrective behaviours to cancel them out. If we take the set ups and results at face value, then they provide a useful gloss on the free-energy minimising underpinnings of retinal stabilisation and the eye tracking behaviours in the oculomotor system.

At the next stage of their experiments the authors wanted to show why actions specified in terms of sensory predictions are optimal. Using the same schematic set up they introduced exogenous interruptions (i.e., unpredictable, hidden effects) to test the behaviours of the system. In this case the process remained undisturbed and was carried out as normal, in other words, the process supporting eye tracking and retinal stabilisation remained insulated from erroneous effects. As such, the authors conclude that active inference facilitates optimal behaviour for:

> […] *if we only have to specify the consequences of an intended or expected*
> *movement, then the actual action will be robust to variations in the true generative*
> *process or unexpected perturbations* to that process. *In other words, specifying*
> *a movement in terms of what we want to see, as opposed to what we want to do*
> automatically makes behaviour more robust (Friston et al, 2010, p. 14)

If we accept this interpretation of the experiments in terms of free-energy minimisation, then what we have is a reasonably clear demonstration of how optimal action can be specified in terms of pre-established expectations about the sensory consequences of movement that are resistant to intervening disturbances.

To develop this line of argument further, Friston used the same experimental set up to simulate active inference during visuo-proprioceptive integration during a visually guided reaching task to provide an account of how visuo-proprioceptive signals are weighted to enhance optimal action. In this case top-down visuo-proprioceptive predictions and bottom-up visual and proprioceptive predictions errors converge to maximise the outcome of multimodal sensorimotor integration during cued reaching movements. In this scenario, a simulated generative model of a two-joint arm which had to reach for a coloured target was used. The proprioceptive input corresponded with the angle of both arm joints and the visual signal corresponded with the position and brightness of the target and the distal extremity of the arm. The same mathematical equations were used to prime the generative model to expect that the distal extremity of the arm would accelerate towards the target only when it was illuminated. This causes the arm joint to move towards the target when it is lit up. In this case the expectation that is driving the movement is visuo-proprioceptive because the generative model is primed to expect that when the stimulus is illuminated the movement of the distal extremity of the arm joint towards the target is both 'seen' and 'felt'.

At the first stage of the experiment, Friston and his colleagues deliberately added a significant amount of noise to the proprioceptive signal, but not the visual signal. What they observed was that the movement trajectory of the arm towards the target was seamless despite the disturbances to the proprioceptive signal. The explanation given for this is that given that if, under active inference, action is elicited to seek out and cancel prediction errors, then the visual and proprioceptive prediction errors will be initially weighted relative to their ability to provide a steady supply of prediction errors. By this, I think they mean errors which make it clear what mistake has been made (e.g., that the hand has not met the target as expected). By adding noise to the proprioceptive signal in such a way that the information provided by the proprioceptive error signal is not clear, but the visual signal remains undisturbed such that it can reliably produce salient prediction errors (errors which the system is more likely to be able to correct for), then the visual signal is the more reliable and stable out of the two and so the visual signal is given more weight. Given that the reaching

movement remained unperturbed by such disruptions we can reasonably assume that this is the case. Thus, we are told:

> This robustness to proprioceptive noise rests on optimising expected precisions in both modalities. The ensuing optimisation boosts visual prediction errors relative to proprioception and ensures that the relative contributions of both modalities is balanced in a Bayes optimal fashion (for perception and action).
> (Friston et al 2010, p.8)

The weighing up of the contributing signals based on their capacity to reliably produce salient prediction errors does two things. Firstly, it allows the motor process to be carried out as normal despite degradation to one of the signals; secondly, it enhances the predictive capacities of one modality relative to the other whenever they have been brought together to produce motor output. To justify this point further, Friston worked out four expected precisions of the visual and proprioceptive signals under high and low levels of noise. What he found was that vision can substitute for proprioception and vice versa, and no disturbance to the movement is observed. Only in cases where both signals were degraded was movement compromised (Friston, 2010, p.10).

This story about why intermodal predictions must be weighed up in a prediction-error minimising, Bayes optimal fashion, is one we're already familiar with for it follows the part of Ernst's MLE principle which states that the weighing up of calibrated signals is, in part, determined by current levels of reliability of each signal. Of course, there are subtle differences over what's meant by 'reliability' here. On Ernst's view the reliability of a signal implies its capacity to provide high-fidelity probabilistic information about whatever property its informational content pertains to. On Friston's view, the reliability of a signal is determined by its capacity to detect relevant prediction errors such that active inference can still take place. Furthermore, the Bayesian view that Ernst adopts is 'bottom-up' heavy in the sense that the informational content of the contributing signals enters the content of the driving signal that determines the outgoing percept. On Friston's view, the bulk of the processing which shapes the outgoing signal is top-down as it is primarily shaped by pre-established cross-modal predictions. In his view, the only way an occurrent incoming signal may have a determinate effect is if it signals that a prediction error has been made.

What is clear under both interpretations is that there is a mounting body of empirical evidence within the Bayesian research paradigm which indicates that multimodal integration is important for maintaining optimal behaviours in the face of sensory signals which vary in their reliability at different times and under different conditions. What's important to both models is that the objective of sensory processing is to maintain a close match between the properties and activities of the body and environment and the sensory signals they generate. Where the prediction error minimisation model gains leverage, I think, is that it has a broader theoretical framework which helps us to make better sense of how these observations about the importance of maximising accuracy/minimising error during cross-modal integration at the sensorimotor level contributes to the optimisation of the activities of the animal when this story is couched in terms of active inference and the free-energy principle. When one considers the constraints placed on the multisensory processes which underpin motor behaviour from the perspective of a framework which helps us better understand how the activities of the agent is constrained at all levels of description, it provides one with a nice rationale for thinking that it is prediction error minimisation which plays an important role in shaping the cross-modal underpinnings of movement given that motoric behaviours must contribute to the optimal behavioural performance of the animal as a whole. In doing so it provides a story about the link between the functional properties of sensorimotor process at one level support the optimisation of motor output at another (i.e., the behavioural level) which Ernst's model lacks[68].

4.9 Challenges for the PEM model of Perception and Action.

Despite the wealth of insight the prediction error minimisation framework offers with respect to the nature of perceptual and motor processes, and their respective and collective functional role, I'd like to finish by laying out some reservations I have with the radical stance that Friston *et al* take with respect to the necessary requirements for motor control. The first worry concerns the radically reductionist take that Friston take on the nature of motor processing and how it supports optimal motor performance. All motor processing consists of,

---

[68] I will also argue that this is one of the defining features of the prediction error minimisation framework that allow us to make headway on the discussions about the significance of the multisensory processes of the body schema for optimising behavioural performance.

on this account, are sensorimotor predictions which elicit motor commands and prediction errors which drive necessary motor adjustments in accordance with the requirements of active inference, or in more formal terms free-energy minimisation. Aside from their explicit statement to this effect that was quoted above, this same line of thought was also observed, for instance, in the commentary quoted above by Friston *et al* (2010) in section 4.6 in relation to their studies on active inference during cross-modal reaching tasks. I refer to their claim that 'specifying movement in terms of what we want to see rather than what we want to do' allows for a seamless movement transition as the motor trajectory remains undisturbed by irrelevant perturbations (Friston *et al* , 2010, p.14). It might be said that this *may* be true at one level of explanation, that is the sensorimotor level, but not true at others, e.g., the intentional level. In other words, even if motor behaviours are produced and guided by sensorimotor predictions at one level, this is only within the context of the higher-level intentions of the agent. For instance, it may be true to say that, at one level, the movement of the arm towards the cup may be produced by visuo-tactile-proprioceptive predictions about the sensory feedback that should be provided when the hand successfully meets the cup and is re-modified in response to incoming cross-modal prediction errors. However, this process takes place within the context of my personal level intention to pick up the cup to take a drink[69]. Andy Clark articulates this worry very nicely when he tells us:

> […] that would not immediately justify us in claiming that it thereby constitutes the cognitive economy. To see this, we need only reflect on the fact that it's all just 'atoms', molecules and the laws of physics too, but that doesn't mean those provide the best constructs and components for the systematic descriptions attempted by cognitive science. The desert landscape theorist thus needs to do more, it seems to me, to demonstrate the explanatory advantages of abandoning more traditional approaches to value, reward and cost[…]. (Clark, 2013, p. 59)

I think this is an important point. Understanding motor production and re-modification in terms of active inference at the level of sensorimotor predictions is important, but this does not, by default, falsify the claim that adaptive intentions and goals play a crucial role in determining and shaping the behavioural strategies biological organisms employ to optimise their interactions in their environment. Still, this worry doesn't undermine the theoretical tools which the AOPP model gives us. Nonetheless, I raise this problem to emphasise that I

---

[69] This line of thought will be picked up again in chapter 6.

do not accept the radical stance Friston and his colleagues adopt. A weaker reading could still allow for the fact that the claims from Friston can be true at one level of explanation and can be incorporated into a model which incorporates the insights from AOPP to explain the link between motor processing and optimal behavioural performance at other levels of explanation[70].

In summary, there is a big concern that currently faces the AOPP model of motor processes that pertains to the overly reductionist take that Friston and co take when they claim that understanding the nature of the motor processes and how it supports optimal behavioural performance can *be purely understood* in terms of the fulfilment of sensorimotor predictions. I argue, following Clark, that whilst this may be true at one level of explanation, there is still scope to incorporate this understanding of the link between the operating principles of sensorimotor integration and optimal behaviour with other levels of explanation, like the intentional level. Thus, I adopt a weaker reading of AOPP which acknowledges that the AOPP story may only be true, but only for one level of explanation[71].

Conclusion

The PEM model of perceptual processing and its extension to motor processing under AOPP hold the promise of transforming our understanding of the nature of the perceptual and sensorimotor system. Both systems are characterised in terms of hierarchical generative models which generate perceptual and sensorimotor predictions to explain incoming sensory data to produce rapid and accurate perceptual and motor outputs respectively. Furthermore, the perceptual and motor system work together under the principle of active inference to minimise, if not cancel out completely, prediction errors. This constitutes an 'action-oriented representation' for maintaining an updated representation of the environment with which to coordinate motor outputs. What is important for our purposes here is the cohesive and seemingly coherent story that the AOPP model provides helps us make better sense of the importance of the underpinning operating principles of perceptual and motor activity, and their role in optimising the behaviours of biological agents in their environment. This model helps us better understand the importance of multimodal integration for prediction error minimisation for motor regulation in real time and motor development over the long term. All

---

[70] This is precisely what I will attempt to do across chapters 5 and 6 of the thesis.

[71] The implications of this point will be explained further in chapter 6.

of this gives us a plausible story about the nature of perceptual and motor processes, their interrelations and how they support optimal performance of biological agents in their environment. As we will see, it gives us the right kind of theoretical tools to begin to build a story about the relationship between the body image and schema and their importance for optimising motor output.

Chapter 5

The Body Schema as Sub-Personal Action-Oriented Representation: A Unifying Account of the Operating Principles of Optimal Multimodal Sensorimotor Processing.

Introduction

The key argument in chapter 2 of this thesis is that a comprehensive analysis of how the body percept and the body schema make independent and collective contributions to the development, transformation and optimisation of motor behaviours requires a clear characterisation of the body schema. Therein we also looked at the respective philosophical analyses of the body schema from Gallagher and Vignemont in terms of a functional description of its components. In this chapter I will use the resources provided by the prediction error minimisation framework to define the nature and operational principles of the body schema in terms of a sub-personal system of hierarchically structured action-oriented representation(s) that is driven by prediction error minimisation, respectively. I will then proceed to illustrate how this characterisation provides a reading of the body schema that is consistent with Vignemont's action-oriented model of the computational structure of the body schema. However, I will argue that my model has more explanatory leverage than Vignemont's and avoids the criticisms I levelled against her model in chapter 2. With respect to Gallagher's model I will argue that my interpretation of the body schema in terms of action-oriented representation is broadly consistent with his functional interpretation of the components of the neurological components of the body schema, and that the active-inference principle helps us make sense of the significance of his functional and behavioural analyses of the body schema. However, I anticipate that this model will invite criticism from Gallagher who argues against representational analyses of the body schema at a later stage of his *2005* where he considers how the body schema needs to be synthesised to best serve the explanatory ambitions of cognition. I conclude that whilst Gallagher's criticism may apply to the sort of representational analysis Vignemont provides, my interpretation of the body schema from the perspective of action-oriented cognitive science can meet these objections and provides a better way of synthesising the insights Gallagher's analysis provides.

5.1. The Philosophical Analyses of the Body Schema: A Reminder of their Motivations and Implications

The pursuit of a story about the significance of the operational principles of optimal cross-modal cue-integration for perceptual and sensorimotor processing is, in part, motivated by the important developments that emerge out of Shaun Gallagher's conceptual dissociation between the body image and the body schema, and his functional/behavioural analysis of the body schema. The gravity of the confusions that have surrounded this concept has been

serious enough for psychologists and one philosopher to argue that the distinction should be dropped altogether. However, it could be said that the confusions don't come about because of the inherently problematic nature of the body schema; rather, the problems emerge out of the failure to respect the conceptual boundaries of the body image and the body schema. In response, Gallagher provides a conceptual distinction between the body image that remains faithful to the traditional characterisation of the distinction. The body image categorises the multisensory, conceptual and affective processes which underpin the personal-level perceptual recognition, understanding and apprehension of the body. In contrast, the body schema categorises the multimodal sensorimotor processes that facilitate and regulate motor behaviours below the threshold of consciousness. As such, both concepts admit of separate questions about how their respective processes fulfil the functional role that is ascribed to each system. The processes of the body image provide a story about how sensory, conceptual and affective processes shape and enable body perception. The processes of the body schema provide a story about the multimodal sensorimotor processes work in conjunction to produce and remodify motor behaviour. As such, Gallagher ascribes to the body schema two components that are responsible for processing and calibrating multisensory information about the position and movement of the body and objects in the surrounding environment, respectively. He also includes a motor component that is responsible for producing motor behaviours by activating all the musculoskeletal networks that are required for the maintenance of balance and motor behaviours. Importantly, his analysis of the interactions between the body schema and the body image shows that body schematic processes do not require personal-level deliberation or monitoring. Recall the example of being visually and proprioceptively aware of your hand reaching towards your coffee cup, but being unaware of your hand forming into the right shape (a grasp) for picking up the cup. Furthermore, the analysis of deafferentation shows that having a body schema that can operate below the threshold of consciousness to produce and regulate motor output optimises our behavioural performance as individuals who are completely reliant on their body percept for the guidance of movement produce motor outputs that are slower, less precise and more cognitively demanding. The important conclusion that was reached with respect to Gallagher's analysis is that we need a story that can take us from his functional analysis of the components of the body schema to a story about why the operations are manifested at the level of behaviour[72].

---

[72] Recall that this analysis comes from Gallagher's *2005*. See sections 2.2, 2.4 and 2.5 in chapter 2 for a reminder.

Recall Vignemont's alternative model of the body schema in terms of an action-oriented representation the processes of which are about, and directed towards, the body[73]. There are three functional components of this action-oriented representation that make up the body schema. The *initial body schema* is involved in the selection of motor output based on prior information about the movement constraints and capabilities of the human body; the *predictive body schema* anticipates the sensory feedback that should be provided by the motor outputs selected for by the initial body schema, and the *updated body schema* provides the actual sensory feedback that's generated by the motor activities of the individual for comparison with the predictions of the predictive body schema. Thus, whatever the body schema is, it involves motor selection that is based on stored information about the motor repertoire of the human body, as well as the interactions between sensorimotor predictions and incoming sensory feedback.

In chapter 2 of this thesis I argued that the conception of the body schema in terms of these three components is the only aspect of Vignemont's model that is plausible and worth preserving. However, there is a caveat for I argued that we should dismiss her account of the sort of action-oriented representation the body schema is insofar as she claims that it is a) conscious and b) its processes only underpin motor behaviours that are directed towards the body. My claim is that this only serves to blur the boundaries between the body percept and the body schema and Gallagher's analysis makes it clear that sub-personal motor processes are directed towards the environment is correct, contra Vignemont[74]. The question is how we can use the theoretical resources from action-oriented cognitive science to make headway on the open issues that the respective functional descriptions of the component parts of the body schema create.

5.2. The Optimal Significance of Multimodal Processing for the Operations of the Body Schema: A Response to Gallagher and Vignemont

With the theoretical resources of the Bayesian models of sensory processing in place, this puts us in a much better position than before from which we can understand the optimal significance of the multimodal integration for the sensorimotor operations of the body

---

[73] De Vignemont: 2010.
[74] See sections 2.8 and 2.9 of chapter 2 for a reminder.

schema. Despite the obvious differences over how the body schema should be conceptualised, the respective models of the body schema from Shaun Gallagher and Frederique de Vignemont emphasise the importance of multisensory processes for optimising the behavioural outputs it produces and regulates. However, this only raises the fundamental question that both interpretations fail to address: *how* and *why* do cross-modal processes optimise motor behaviour. At this stage, we can answer the 'how' question in the following way. In accordance with the general thread of the Bayesian models we can say that cross-modal inputs are brought together and weighed up based on their context and task-specific reliability to maximise the accuracy of the driving signal they support. The flip side of the coin is that the maximisation of accuracy minimises the possibility of error. Furthermore, having multiple sources of sensory information that carry equivalent information means that, in the event one malfunctions or is further degraded, another can override it so that the process can carry on as normal. Thus, the answer to the 'why' question from Ernst and Friston is that having multiple sources of concurrent sensory information from the distinct modalities which are brought together to maximise the accuracy of the signal (by minimising errors) makes for a more stable, robust and efficient process. For Ernst, the answer to the 'why' question stops at why error minimisation optimises the output. However, Friston & co provide the additional story that optimising every process in real time equates to minimising free-energy over the long term and, in doing so, optimises the interactions of biological agents in their environment. The next question is, how do we interpret the components of the body schema from the perspective of AOPP that is consistent with the standard definition? I address this question in the next section.

5.3. An Analysis of the Body Schema in Terms of a Prediction Error Minimising, Sub-Personal, Hierarchical Action-Oriented Representation

The most plausible functional description of the body schema conceives of its components in terms of multimodal sensorimotor processes that calibrate cross-modal inputs and produce motor output below the threshold of consciousness. So defined, the prediction error minimisation framework as it is applies to action, i.e., action-oriented predictive processing (or AOPP as I refer to it in chapter 5) would conceive of this system in terms of a sub-personal hierarchical action-oriented representation (or set of representations) that elicits

cross-modal sensory predictions about the state of the body and environment and sensorimotor predictions about the sensory consequences of its movements. The sensory predictions pertain to the probabilistic causal relationship between the sensory apparatus and the state of the body and environment (i.e., the position and movement of its body and objects in the environment). In turn, the sensorimotor predictions pertain to the internal changes (i.e., the sensory stimulus that should be fed back to the system) that come about as a result of the individual's motor behaviours. As such, the range of motor responses produced by the body schema would, in part, be shaped by a recognition ensemble in the brain which defines the permissible (and impermissible) parameters of action for the individual based on the internal state changes those actions bring about, i.e., whether they bring about causal changes to the internal states of the agent that maintain its homeostasis (or not). Furthermore, the body schema includes sub-personal cross-modal predictions about the current motor state of the body (i.e., its postural position and movement) and the state of the environment which constrain the range of available actions that are available to the agent. Furthermore, sensory inputs provide the necessary feedback required for verifying the veracity of the predictions and, more importantly, for the identification and minimisation of prediction errors (i.e., deviations from anticipated sensory feedback). So conceived, the body schema is a hierarchically structured, sub-personal, action-oriented representation that is shaped by prediction error minimisation.

To get a clearer sense of how this characterisation of the body schema works consider again the example of your hand forming into a grasp to pick up your coffee cup outwith the scope of your awareness. Your recognition ensemble enables and restricts the range of possible actions that you can take with respect to the mug because there are more and less effective ways of interacting with the cup; you can grasp it with two hands or with only one hand, you can cusp your fingers around the surface of the mug, or you can grasp the rime of the mug with the tips of your fingers, and so forth. Therefore, there are many ways you can grasp the mug. However, the available range of actions might be narrowed down further on the basis of probabilistic sensory predictions about relevant aspects of the environment or relevant contextual cues; for example, the three-dimensional configuration of the cup, such as its size, might determine how you grasp the mug. Grasping the handle of a particularly large mug with your fingers might be more effective if your hands are too small to grasp the body of the mug firmly. Furthermore, if you are reaching your hand down towards my coffee cup when you are driving your car it is more effective to grasp the mug with a firmer grip compared to the kind of grip you might have on the mug when you pick it up to clear it away

(I'll come back to this point in due course). Therefore, the sub-personal top-down processes of the body schema are important for defining and re-defining the selection of the motor outputs for they will play a key role in determining the sensorimotor predictions that will be elicited to produce the selected motor behaviour (e.g., corresponding proprioceptive and tactile feedback which report that the fingers have latched onto the handle of the cup in the intended way). In turn, the incoming sensory data the motor trajectory generates provides necessary feedback to determine whether the movement is unfolding in the anticipated way, or not.

In sum, a description of the functional components of the body schema in terms of a sub-personal action-oriented representation in accordance with AOPP goes as follows. The body schema is part of the individual's internal generative model which elicits probabilistic predictions about the state of its environment and the two-way causal effects of its behaviours on its embodied states and its environment, and vice versa. More precisely, the recognition density defines the motor repertoire of the body schema in accordance with the free-energy minimisation principle and, together with sub-personal sensory predictions about the current state and activities of the body and the environment (as specified above), selects the appropriate sensorimotor predictions that produces the selected motor behaviour. Any relevant prediction errors are quashed by eliciting alternative, or additional, sensorimotor predictions to bring about necessary changes to the motor outputs in accordance with the active inference principle[75].

5.3. A Response to De Vignemont

This story about the kind of system the body schema is (i.e., a hierarchical multimodal sensorimotor system) is compatible with Vignemont's functional description of the different aspects of the components that make up her action-oriented interpretation of the body schema in some ways, but not in others. Let me explain. On her account, the *initial body schema* includes prior information about the motor repertoire of the human body and is involved in the selection of motor output based on their capacity to achieve the motor outcome in the most economical way. The model I provide here offers the same proposal but also explains

---

[75] This definition suffices for the present purposes of this section.

how, and why, the parameters of the motor repertoire are so determined. From the perspective of the free-energy minimisation principle the more economical an action is, the more it minimises free-energy because it uses fewer of the agent's predictive resources and ensures the best fit between the sensorimotor predictions of the individual's generative model and the sensory evidence, in comparison to other available actions with the same functional properties. Minimising free-energy, in the context of action, is an optimising strategy for the selection of motor behaviours which produce outcomes that support the activities of the individual, make fewer demands on its cognitive resources and ensure that equilibrium is maintained within its internal economy contributes to the long-term survival of biological agents over the long-term.

This brings me to the next aspect of the body schema - the *predictive body schema* which elicits sensorimotor predictions about the sensory consequences of the selected movement, i.e., motor induced sensory feedback that is exclusively about the body. The predictive component of my interpretation of the body schema cuts across the sensory predictions about the state of the body and the environment, and sensorimotor predictions about the movement of the body, for reasons I have elucidated in my interpretation of the example above. Vignemont's *updated body schema* provides new and updated sensory feedback for comparison with the predictions of the predictive body schema. My interpretation of the body schema also includes this component but has additional explanatory leverage over her account because it provides the theoretical resources to explain the importance of the incoming sensory data for the overall operation of the components of the body schema. The incoming sensory data verifies whether the predictions have been fulfilled in the anticipated way, and prediction errors drive a change to the predictions to bring about a change in motor output.

The overarching implication of this pertains to the criticisms I levelled against Vignemont's model in chapter 2 pertaining to her objection to Gallagher's account of the body schema and her account of the kind of action-oriented representation the body schema is. First, I argued that, contra Vignemont, simply providing a functional description of the components of the body schema in terms of internal representation does not give her model any additional explanatory leverage until she provides the additional story of how the neurophysiological underpinnings of the body schema carry out the functions ascribed to them. In comparison, by drawing upon the theoretical resources of AOPP to mount my interpretation of the body schema, my model can provide such a story in the precise way I have outlined here. Furthermore, I argued that her account of action-oriented representation is problematic on

several grounds. First, it goes against the grain of standard conceptions of action-oriented representation; second, there is no rationale for thinking that the processes of the body schema are exclusive to the behaviours that are directed towards the body; third, her account of the body schema as a conscious action-oriented representation not only perpetuates the conceptual confusion surrounding the body schema, but her line of justification is based on an inconsistent analysis of the body percept and the body schema.

My interpretation of the body schema avoids these problems for it is consistent with the standard conception of action-oriented representation. The key difference, in light of the considerations from chapter 4, is that the sensorimotor processes that underpin the representation(s) of the body schema are probabilistic predictions. Furthermore, through my analysis of the example of grasping the hand, I have made it clear that such an action-oriented representation is sub-personal and is an open feedback loop that is constantly engaged with the environment. Therefore, to the best of my understanding, there are no inconsistencies between my interpretation of the body schema in terms of action-oriented representation and the usual requirements for a satisfactory conception of the body schema[76].

5.4. An Explanation of how the Prediction Error Minimising Processes of the Body Schema Optimise Motor Behaviour: The First Reply to Gallagher.

Likewise, the model of the body schema I offer here is consistent with Shaun Gallagher's interpretation of the functional structure of the underlying neurological components of the body schema. The first component registers cross-modal sensory information from the low-level sensory systems and sends them to the 'intermodal capacities' which 'facilitate the communication' between cross-modal inputs. According to my interpretation, this process is a top-down process where the communication between cross-modal signals entails the temporary coupling of the cross-modal signals to provide task-specific, contextually appropriate, probabilistic multisensory predictions about the state of the environment and the body. The communication also occurs in a bottom-up fashion by way of prediction errors. Thus, under my interpretation the communication between the sensory processes of the body

---

[76] The interpretation of the table is also consistent with the initial characterisation of the body schema in terms of a 'combined standard, against which all subsequent changes of posture are measured' (Head & Holmes, 1911, p.187). The 'combined standard', on my interpretation, is the cross-modal sensorimotor predictions.

schema are top-down and bottom-up. Furthermore, top-down motor processes are prior information and sensorimotor predictions. To this extent, the models are not incompatible for the key difference boils down to a difference over the shape of the underlying information processing.

The model I offer also has additional explanatory advantage insofar as it puts us in a much better position to make better sense of Gallagher's observations of the relationship between the operations of the body schema at the functional level and how their workings are manifested at the behavioural level. If the body schema can generate sensory and sensorimotor predictions about the probable current state of the body and the environment, and of quashing incoming prediction errors, independent of the body percept, then motor responses can be produced and regulated without making so much demands on our cognitive resources. From the perspective of free-energy minimisation, it makes good adaptive sense to have a system that can facilitate the same tasks as the body percept, but which utilises fewer of our predictive resources. The case of Ian Waterman shows that behaviours that require, and are produced by, continuous personal-level perceptual guidance are slower, less precise and place more demands on our cognitive resources[77].

This raises the important question of how the relationship between the body image and body schema gets to this point. For Gallagher, a conceptually satisfying account of the body schema can't just make reference to its neurological components for:

> Without a certain amount of selectivity, *built up by practice and the cultivation of habitual movements,* the body might move in any one of multiple ways, since *the possibilities allowed by physiology are much greater than the particular movements necessary…Thus the body schema is much more selectively attuned to its environment than what physiology on its own will specify* (Gallagher, 2005, p. 143, emphasis mine.)

Here Gallagher is expressing the point that the interactions between the body and environment over different timescales put further constraints on the outputs of the body schema, i.e., over the long-term and in real time. When he says that 'the body might move in

---

[77] This is much as I can say at this stage about how the body functional operations of the body schema are manifested at the behavioural level. A full appreciation of this line of thought can only be provided by an interpretation of Gallagher's functional and behavioural analysis of the body schema and the body percept from the perspective of the predictive processing framework. Thus, I will develop this point in the next chapter.

any of multiple ways, since the possibilities allowed by physiology are much greater than the particular movements necessary' he is expressing the same point I mentioned earlier about the fact that the constraints set by the body on action are not sufficient to narrow the selection down for the current state of the environment will put further constraints on action. Hence, we arrive at the claim that the body schema is 'much more selectively attuned to the environment than what physiology on its own will specify'. The overall point Gallagher expresses is that the body schema is a 'holistic system'[78] wherein 'the brain attunes itself to what the body and environment affords'[79]

It is quite difficult to provide a fully fleshed story without making reference to the body image in light of Gallagher's analysis of how the motor functions of the body schema are built up. Recall his analysis in chapter 2 pertaining to the role of the body percept (i.e., the cross-modal processes which underpin perceptual awareness of the body) in shaping motor behaviours that are then deferred to the body schema once they have been sufficiently mastered. The prediction error minimisation is a story about how we come into meaningful perceptual and motoric contact with our body and environment. In other words, how we become sensitive and responsive to our body and environment. Successful perceptions and actions are brought about when equilibrium is met between the predictions of our generative model and the sensory feedback our perceptual and motor processes generate. This process is built up through a process of active inference in real time and is regulated over the long-term by free-energy minimisation.

At the point in time when the body schema is producing behaviours without the intervention of the body percept, it is because it has undergone an extensive prediction error minimising period where the motor behaviours of the body have been monitored at the personal-level to test out the sensory and sensorimotor consequences of the behavioural routines in question. Once the prediction errors have been minimised to such a degree that the individual can carry out the process without perceptual monitoring, the sensorimotor predictions are then subsumed to body schema to regulate the behaviours below the threshold of consciousness. Every time the body schema elicits motoric interactions it provides another opportunity to test out the accuracy of the sub-personal models of the body schema about the causal interactions of the body in the environment. Thus, is it not the case that action-oriented representations are somehow closed off from the surrounding environment for they do not:

---

[78] Gallagher, 2005, p.143
[79] Gallagher, 2014, p.100

aim to *engage* the world in some action-neutral fashion, and they are firmly rooted in the history of organism-environment interactions the sensory stimulations that installed the probabilistic generative model. (Clark, 2015, p. 4)

Remember as per the free-energy principle, different organisms with different embodied configurations and homeostatic requirements will possess their respective species specific generative models that are further customised along developmental and life-long trajectories. The upshot is that even with its emphasis on the internal models and representations, AOPP does not overlook the fact that the particular embodiment an animal has, and its environmental interactions, are crucial to the nature and functioning of action-oriented representations (Clark, 2015, pp.3-4).

5.5. The Relationship between the Body Schema and Intentions: The Second Response to Gallagher.

A related point is that as much as the agent's particular embodied configuration and its environment play an important role in enabling and constricting its general motor repertoire, this can't be the whole story for an individual's responsiveness to their surroundings is often shaped by their intentions for action, and this narrows down motor selection further. The body schema, Gallagher tells us, is not an intention; rather, it makes an important contribution to the intentional purposes of the agent[80]. Studies from Jeannerod *et al*[81] lend empirical support to this claim. They showed that the rotation of the wrist is different depending on whether an individual intends to pick up an object or throw it away. The question, then, is what sort of contribution the body schema makes to the fulfilment of intentions. The place to start is by providing an explanation of what we mean by intention. Clearly, the working notion of intention at play entails the intention to do something. More particularly, it involves the personal-level intentions of the agent, such as intending to pick up the cup, throw the ball, open a door, etc. Thus, the issue in need of explanation is how the processes of the body schema work to fulfil the intentional purposes of the agent.

---

[80] Gallagher: 2005, pp.22-26.
[81] Jeannerod, Jacob: 2003.

As a first step, I suggest that we adopt Ruth Millikan's conception of personal-level intentions in terms of a pushmi-pullyu representation (PPRs). Pushmi-pullyu representations have 'descriptive content'- information about the current state of the world is - and directive content which is a motor instruction (or class of motor instructions). Here is her description of intentions in terms of PPRs

> Human intentions are probably an example of PPRs in thought, serving at once to direct action and to describe ones future so that one can plan around it. (Millikan, 1984, pg.1)

Thus defined, PPRs that underpin human intentions produce motor behaviours on the basis of the anticipated outcomes they bring about for the agent. For all intents and purposes, such PPRs are a species of personal-level action-oriented representation for they direct motor outputs based on their capacity to bring about the intended and/or expected outcome they are recruited to produce. The question is how the processes of the body schema are recruited by, and assist the activities of, this kind of PPR. To explain, I will appeal to Susan Hurley's[82] distinction between basic and non-basic intentions. A non-basic intention is the intention to perform a action, like picking up an apple or pushing a button. Basic intentions instruct the non-conscious movement of the body; in other words, a basic intention is a motor intention such as moving the finger or shifting my gaze that is performed without my knowing I am doing this[83]. According to Hurley, basic intentions are 'where your intentions begin' (I'll come back to this point)[84].

The question is how this bears on our interest in explaining the responsiveness of the body schema to the personal-level intentional activities of the agent. In accordance with the definition between the body image and the body schema, non-basic intentions would include personal level intentions to move my body in the way required for performing the intended action (like picking up the coffee cup). Thus, non-basic intentions are the province of the body image. Basic intentions (such as forming the hand into a grasp) are the province of the body schema. Thus, to tell a story about how the processes of the body schema are, in Gallagher's words, 'attuned' to the intentional purposes of the individual is to tell a story

---

[82] Hurley:1998.

[84] Hurley: 1998, p.357

about how basic intentions are recruited by non-basic intentions. Hurley's line of thought is that the development and transformation of non-basic intentions is built up through a process of using pre-existing motor strategies (i.e., other non-basic intentions and basic intentions). Crucially, the selected pre-existing motor strategies are chosen based on their predictable outcomes:

> Sensory feedback allows me to select among already available intentional actions one that has the *desired* effects, thereby bringing those effects under intentional control. This is similar to what happens when we discover that something we can do has a *reliable consequence* we were previously aware of. We thereby acquire a new description under which such acts can be intentional (Hurley, 1998, p. 366).

A complete story of how the basic intentions of the body schema are developed and transformed can only be told through an analysis of the interaction between the body percept and the body schema over the course of development. Such a story goes beyond the scope of this chapter[85]. However, we can still use Hurley's line of thought to provide a rough-and-ready sketch about how the operations of the body schema are responsive to the intentional purposes of an individual in real time for her account meshes well with AOPP.

To begin I invite the reader to cast their eye back to Clark's point that I outlined in chapter 4[86] where he says that in most normal circumstances a lot of the representational activities of the agent will have already been brought into equilibrium as their sequences of motor behaviours are produced. Therefore, to intend to pick up the cup is to intend to bring about a new series of visual and proprioceptive sensory feedback. Thus, to intend to pick up the cup is to elicit predictions about its sensory profile (e.g., visual predictions about its size) and predict that picking up the cup requires the generation of a certain sequence of, say, visual and proprioceptive feedback. Thus, when I intend to pick up the cup I elicit a sequence of body movements that include the non-basic intention to pick up the cup and, in turn this elicits relevant basic intentions, e.g., the formation of the hand into a grasp. This non-basic intention is constituted by personal-level predictions (i.e., about what I should see and feel as my arm stretches over the cup) and sub-personal predictions (e.g., the visuo-proprioceptive predictions about the feedback that should be provided by the hand forming into a grasp and coming into contact with the cup.) By this point, after a long history of picking up objects of

---

[85] In the next chapter I will use Gallagher's analysis of the interactions between the body schema and the body percept to make better sense of Hurley's story.
[86] See section 4.5

its size, the grasp will have been incorporated into this motor sequence on the basis that it has successfully fulfilled the underlying visual-proprioceptive predictions of similar intentions in relation to objects with that sensory profile before.

5.6. Why an Analysis of the Body Schema in Terms of a Prediction Error Minimising, Action-Oriented Representation is *not* just Another Instance of 'Body Snatching': The Third Response to Gallagher

I provide the aforementioned responses to Gallagher in anticipation of a line of objection that I think he might mount against my position. His claim that the operations of the body schema cannot be fully understood and explained in terms of its underlying neurological components on account of its sensitivity to bodily and environmental constraints and intentions is part of a much bigger argument he levels against representational analysis of the body schema in his *2005*. In turn, this reflects his anti-representational stance about mental and cognitive phenomena in general.

   To set the scene, consider the following point. The intended characterisation of the body schema and body image is ripe for an analysis from the perspective of embodied cognition. Why? Well, the body schema is supposed to explain how the central system and peripheral sensory systems work together to localise the position of the body parts in space. The body image is supposed to help explain how information from the peripheral sensory apparatus is processed for the purpose of producing perceptual awareness of the body. This is just another way of saying that the body schema and body image can help us to analyse how low-level processes of the body (i.e., the sensory apparatus) works in conjunction with the brain to produce motor output and body perception, respectively. In chapter 6 of his book, Gallagher makes it clear that a proper conception of the distinction can take us some of the way towards a theory of embodied cognition, but not all the way. For example, an analysis of the body image can help us explain how the body appears in the perceptual field, whereas the body schema can help us understand the processes that occur 'behind the scenes' to constrain the perceptual field[87]. However, a complete story of human cognition requires the synthesis of studies from neurology, behavioural studies to studies in phenomenology.

---

[87] Gallagher, 2005, p.138

With that said, the first step Gallagher takes is to criticise some of the major contenders in in philosophy of mind and cognition which have banked on the assumption that a full explanation of cognition doesn't need to consider the contribution made by the body. Cartesian dualism is one example he cites where the role of the body is, one the one hand, completely denied, but afforded dubious status at another stage of the argument. In his Meditations Descartes states that the mind can be analysed independently of the body and, in another, acknowledges the intimate link between the mind and the body[88]. As part of his ontological and metaphysical arguments for mind-body dualism, he states, roughly speaking, that the mind and body are two ontologically distinct substances in the sense that the former is immaterial and the latter material (hence, we arrive at 'Cartesian substance dualism'). As such, it is logically possible that each can exist independently of the other. One important implication, amongst many, of Cartesian substance dualism is that a full understanding of the nature of the mind can be secured via an explanation of its ontological, metaphysical and epistemic properties. Broadly construed, this amounts to the idea that the mind is an immaterial substance which can exist independently of the body, and can only be accessed and known through introspection, respectively[89]. Of course, in paragraph 20 of Meditation VI an interesting tension emerges with respect to the relation between mind and body: the 'mind-body problem'. One version of the problem[90] is how to explain how the material body can causally influence an immaterial mind during our day-to-day lives, and vice versa. For example, bodily events such as the rubbing of the linings of the stomach and a decrease in fluids, cause mental events like the experience of hunger and thirst, respectively. In turn this fosters the desire to satiate our hunger and thirst (another mental event) which causes us to move our body in the direction of sources of food and water (another bodily event). The apparent site of the interaction, the story goes, is the brain because 'the mind is not immediately affected by all parts of the body, but only by the brain' Descartes, 1641, Med.6. par.20). Therefore, even though Descartes provides a clear acknowledgement that the body and mind are inextricably bound up in intimate ways, insofar as the body does influence the mind it does so only indirectly via the effects of its activities on the brain.

A more contentious approach for Gallagher is one which purports to acknowledge the importance of the (non-neural) body, but fleshes out its contribution in terms of how the body is represented in the brain. This is a problem that Gallagher has humorously called the

---

[89] See Descartes (1641), Meditation II.
[90] The other is to explain how, as Descartes claims, 'my mind and body compose a certain unity' (Meditation VI, par. 20)

'[91]invasion of the body snatchers'. Body snatchers are those who have infiltrated the embodied cognition camp and who come guised as embodied cognition theorists, but are covert computationalists and/or representationalists. The purported problem with such theorists is that they presuppose that all that matters for a complete analysis of the nature and operation of mental and cognitive phenomena is an explanation of the underlying mental computations and, in turn, an analysis of the internal representations which mediate the inputs and outputs (i.e., what is happening in the brain. As such, body snatchers are interested in the contribution of the body to mental and cognitive processes only insofar as it contributes to the operations of the brain, but even then, its importance is only considered to be instrumental. This approach is what Hurley[92] called the 'Sandwich' conception of cognition where perceptual inputs (one slice), cognition (the meat) and motor outputs (the other slice of bread) are causally related in a linear fashion, and cognition is the only real source of interest. Insofar as perception and action are important, it is only because they provide the means of facilitating 'real' cognitive processes.

Therefore, Gallagher claims, representational accounts of the body schema (and cognition in general) are doomed to fail. In virtue of their exclusive focus on the operations of the brain they can only provide an incomplete analysis of what the operations of the body schema are doing after sensory inputs arrive from the periphery and motor outputs are selected. What they don't consider is the influence that the body has on the operations of the body schema, nor the environment. A good illustration of this is Vignemont's account of the body schema as a system of processes which exclusively direct the movements that are directed towards the body. I argued in chapter two that she is mistaken in her assumption that an account of the body schema in terms of the functional structure of its neurological components is explanatorily superior to Gallagher's non-representational account. Vignemont supposed that by conceptualising the functional components of the body schema in terms of representations gave her account additional explanatory leverage over Gallagher's model. The problem with this, I argued, is that it simply does not without an additional story about how the neurological components fulfil the function they are ascribed. Furthermore, her conception of the body schema only served to fuel the confusions. Even more serious is the fact that representational accounts of the body schema have, up to date, only led to a dead end.

However, here is an important point. In my view, Gallagher's position faces the same problems insofar as he only provides a functional description of the components of the body

---

[91] Gallagher:2015.
[92] Hurley:1998.

schema. Extending this account to say that the components are sensitive and responsive to constraints set by the body, the environment and intentions is one thing, but providing an explanation of how this comes about is another. This is where I think my account has explanatory leverage over both accounts of the body schema for it *does* have the underlying story about how the functional components of the body schema fulfil their function along different timelines. Andy Clark expresses this point beautifully when he says:

> […] *it is surely that very model invoking schema that allows us to understand how it is that these looping dynamical regimes arise and enable such spectacular results.* The regimes arise and succeed because the system self-organises around prediction error so as to capture organism-salient patterns at various scales of space and time in the partially self-created input stream. These patterns specify complex, inter-animated structures of bodily and worldly causes. *Subtract this guiding vision and what remains is just the picture of complex looping dynamics spanning brain, body and world.* (Clark, 2015, p.6 emphasis mine)

The point is that without this story, all we would be left with is the acknowledgement that the body and the environment are important for enabling and constraining our myriad of mental and cognitive processes. Furthermore, much like the similar confusions that surround the concept of the body schema, the use of the term representation is not, in principle, problematic. The problem is how it has been used to interpret the body schema. The real difference between representational and non-representational accounts, it seems to me, is a difference in emphasis about the scope of the operations of the body schema. However, I hope to have shown here that the account I offer can help work towards bridging the gap between the opposing models with its emphasis on the operating principles which determine the body schema across development and in real time, and the accompanying explanation of how the body schema builds up its responsiveness to the environment and intentional purposes of the agent to optimise its behavioural interactions.

5.10 Taking stock

The principal argument of this chapter is that thinking of the functional multimodal sensorimotor components of the body schema as constituting a sub-personal action-oriented representation(s) which are shaped and driven by active inference/free energy minimisation can help us make good progress on understanding the optimal significance of multimodal processing for motor control. It also helps us make better sense of the pre-existing interpretations of the body schema from Gallagher and Vignemont. With respect to Vignemont's model I stressed that my interpretation is consistent with hers in the sense that she outlines the body schema in terms of an action-oriented representation that consists of priors about the movement constraints of the body, sensorimotor predictions and sensory feedback. However, I offer a more plausible reading of the body schema in terms of action-oriented representation in comparison to Vignemont for on my interpretation the body schema is *not conscious* and its operations are not exclusively directed towards the body. Moreover, where her model fails to provide a causal explanation of how the functional components of the body schema fulfil their function, I offered here the additional background story about the operating principle of sensorimotor processes from the AOPP framework: prediction error minimisation. With respect to Gallagher's model I have argued that the interpretation on offer is consistent with his broadly functional assessment of its neurological components and can go some way towards meeting the requirements he sets for a comprehensive analysis of the body schema. In doing so, I have attempted to offset a potential objection from Gallagher on the grounds that my interpretation analyses the body schema in terms of internal representation by explaining how it is not as neurocentric and reductive as pre-existing accounts of the body schema have tended to be. Furthermore, I claim that without the necessary story about the underlying principles at play, Gallagher's model has no explanatory advantage over Vignemont's model until his account is supplemented with a story about the processes which drive and optimise the processes of the body schema. As such, I hope to have shown how the model I provide straddles both sides of the divide in a more productive way. However, this is only part of our story for providing a definition of the body schema and providing an explanation of the nature of its multimodal sensorimotor processes is the first step towards a theory of the body image and body schema distinction. Therefore, our next step is to reconsider the Gallagher's behavioural analysis of the interactions between the body percept and the body schema from the perspective of PEM

111

and AOPP to expand this story about how multimodal processes optimise perceptual and motor outputs, and address important deficits in developmental psychology. The next and final chapter, which will bring to bear the insights of the action-oriented model on Gallagher's behavioural analysis of the body image and schema distinction, will act as a response to this challenge.

Chapter 6

A Re-evaluation of Gallagher's Behavioural Analysis of the Body Image and Body Schema from the Perspective of the Prediction Error Minimisation Framework: Towards a Comprehensive Theory of the Relationship between the Body Image and the Body Schema.

Introduction

My purposes with respect to the body schema and body image distinction are twofold: First, to make good progress on our understanding of the body schema, and how its processes optimise behavioural output; second, to build upon Gallagher's behavioural analysis of the body percept and the body schema. Having focused primarily on the issues pertaining to the body schema in the preceding chapters, in this final chapter I will re-examine Gallagher's interpretation of the body percept and his analysis of the role it plays in the optimisation of behavioural output under conditions in which an individual whose sensorimotor skills are under-developed must use their body to successfully perform an action, or must use their pre-established motor skills to successfully navigate their environment in conditions that could have potentially risky or fatal consequences. The overarching conclusion is that the predictive processing framework that has been advanced in this thesis can help us make better sense of Gallagher's line of thought. The common feature of both cases is that they require an individual to act in conditions that present the risk of making substantial motor errors and the central claim of this chapter is that the best way to optimise behavioural performance under such conditions is by way of perceptual and active inference. The broader implication is that this puts us in a much better position to explain the underlying operating principles at play in the case studies that were discussed in chapter one of this thesis. Moreover, it gives us the right sort of theoretical resources to begin the proceedings of working towards a Bayesian style neuroconstructivist analysis of the interaction between the body image and body schema. As such, it allows us to address two important deficits in developmental psychology: the lack of a theory of the body image and body schema and the lack of a causal explanation of how the body image and schema interact across different timelines.

6.1. The Divisions over the Future of the Body Image and the Body Schema Distinction in Cognitive Science.

As the line of discussion in chapters one and two of the thesis unfolded it became increasingly more clear that the future of the distinction between the body image and the body schema is in jeopardy because of the conceptual confusions that have loomed over the latter. Some support for this position is given by the fact that more contemporary definitions of the concept of the body schema still exhibit inconsistencies over the issue of where the line

between the sub-personal body schema and the personal level body image should be drawn up. For instance, recall the definitions from chapter one which conceptualise the body schema in terms of the *awareness* of the spatial position of the body and/or the *understanding* of its movement constraints[93]. Such characterisations implicate the body schema in personal level awareness and understanding when it should not be; rather, such capacities are the province of the body image. Nonetheless, I argued against the approaches taken towards disposing of the distinction on three grounds. First, those who argue against the retention of the distinction misdiagnose the source of the problematic conclusions their analyses of the body schema give rise to; second, the alternative proposal to replace the distinction from Charles Spence & Nicholas Holmes has never materialised; third, the attempt to redefine the body schema from Vignemont before she concedes that the distinction should be jettisoned muddles the boundaries of the distinction between the body image and body schema. This is the upshot of chapter two of the thesis. The overarching conclusion of these preliminary chapters is that the situation that emerges from these arguments is one in which prospective researchers are left with multiple, and unworkable, definitions of the body schema and alternative proposals which either muddy the confusion further or are half-baked at best. Either way, the current state of the debate seems to have reached a stalemate for those who were interested in providing analyses of the body image and body schema distinction have all but given up and a viable alternative explanatory model is yet to be established. I have taken the lead from Shaun Gallagher and re-defined the concept of the body schema in terms of cascading forward models that are part of a much bigger hierarchical action-oriented representation. Up to this point, I have demonstrated how this redefinition helps address the key points of difference in the debate between Gallagher and Vignemont over how the body schema should be conceptualised and the explanatory import of the concept. Now I want to extend this approach to the important points raised in chapter one of the thesis about the place of the distinction between the body schema and the body image in perceptual psychology for another important step towards rectifying the current situation is to re-evaluate the distinction considering the current interpretation of multisensory integration on offer. More specifically, the emphasis of this chapter will be on very recent developments in developmental psychology.

---

[93] See section 1.6 for a reminder.

6.2. How to make a Proper Distinction between the Body Image and the Body Schema

The emerging question is how to apply the insights gleaned from the model of perceptual and motor processes that has been advanced in this thesis under the predictive processing scheme to a theory of the body schema, the body image and their interrelations. The approach I will take up will involve laying out the distinction as it would be carved up under the PEM and AOPP models of perception and motor processes respectively, and a re-examination of the important studies from chapter one and two about how motor-generated cross-modal relations shape and transform the perceptual apprehension and understanding of the perceptual features and motor capabilities of the body. In short, I will explore the role of active inference in determining the relations between motor control and body perception.

Differences in the sensory information that the body image and the body schema work over is nicely demonstrated in the case of Ian Waterman. Prior to the neuropathy which destroyed his motor capabilities his behaviours were guided by visual and interoceptive proprioception, i.e., *sub-personal* information about the postural configuration of his body that is generated by proprioceptors in the muscles, tendons and joints. Nonetheless, there is another source of sensory information that provides the same information that *can* directly structure the body percept: visual proprioceptive awareness. It is this latter source of information that Ian relies on to produce and regulate his motor behaviours after his neuropathy. In Ian's case, visual proprioceptive awareness that is provided by his (still intact) visual body percept has compensated for his lack of visual and interoceptive proprioception that was usually produced by his (now defunct) body schema. If we construe this under the predictive processing framework, then it amounts to the idea that the body image and schema are distinct insofar as they consist of distinct generative models that dispense different predictions, some of which scale up to the level of personal awareness (i.e., the sensory processes of the body image – the body percept) and some of which do not (i.e., the sensory processes of the body schema).

6.3. Current Issues Pertaining to Body Perception and Motor Control in Developmental Psychology

In chapter one of the thesis I gestured towards the idea that a new window of opportunity has opened for the distinction between the body image and body schema to do useful explanatory work within more contemporary developments on the interrelations between motor control and body perception in developmental psychology. An emerging body of empirical evidence suggests that a primordial body schema and body percept develop prenatally and continue to develop and transform across the course of development (more on this below). It is unclear, and indeed seems hard to determine, when the body schema could be considered 'complete'. We continually undergo motor transformations throughout our entire lifetime as we pick up new skills and lose our capacity to perform certain behaviours, say through neuropathy or lack of practice over an extended period. However, the evidence would suggest that developing infants have a complete body percept by the age of two at which stage he or she displays an appreciation of his or her body not as just the aggregation of distinct body parts, but as a unified whole[94]. The principle matter of interest to developmental psychologists when it comes to the body schema and the body image is the causal influence that each has on the other, if any. This is with a view to explaining how the perceptual apprehension of the body, motor control and their interrelations develop and are transformed out of sensorimotor couplings. For Manos Tsakiris[95] these issues represent an important explanatory gap in developmental psychology that only an understanding of the interrelations between the body schema and body image can address. He bemoans the fact that more ink has been spilled on the issue of how to separate these two concepts relative to the issue of how the different systems that fall under each concept work together. The reason for thinking that a story about the interplay between the body schema and the body image is important comes from a growing body of observational evidence which would suggest that developing infants' perceptual understanding of their body is important for his or her grasp of what she or he can do with it. Consider this extract from Holmes & Spence about a 4-month-old baby who, after a period of curious visual fixation with its hand, exhibited surprise and distress when it failed to grasp its own hand because he or she fails to appreciate the fact that whilst one can grasp one's hand using the opposite hand, say by using the left hand to grasp the right hand, one cannot, for instance, use one's right hand to grasp one's right hand. Here is an account of what was observed:

---

[94] Brownell, Svetlova & Nichols in Virginia & Slaughter, eds: 2014.
[95] Tsakiris in Virginia & Slaughter, eds: 2014, p.23-26.

> Sometimes the hand would be stared at steadily, perhaps with growing intensity until intensity had reached such a pitch that a grasping movement followed as if the infant tried by an automatic action of the motor hand to grasp the visual hand and it was switched out of out of the centre of vision and lost as if it had magically vanished (Hall, 1898, p. 351) in Holmes, Lewkowicz & Spence, 2012, p.113)

The important upshot is that the surprise exhibited by the infant demonstrates a failure to appreciate the movement constraints of its body and this feeds directly into its failure to select a more appropriate action. As such, it represents a case where a failure to understand the motor potential of the body at this young age leads to motor errors.

Scale errors are a manifestation of a similar phenomenon, only in these cases the developing infants' perceptual misapprehension of the scale of their body, its size in this case, is the cause of a related motor error. To get the picture consider the following. Mature adults in possession of a healthy and normally functioning perceptual system and high-level cognitive faculties understand that the size of their body matters when it comes to manoeuvring their body in their environment. For example, the height of an individual will determine whether he or she can walk through an entrance without having to duck their head, and the breadth of their body will determine whether he or she can fit through narrow spaces. For the purposes of the ensuing discussion it is important to acknowledge the fact that the body of an adult human is much too large to fit in a small car that is no bigger than the size of their foot. In contrast, young toddlers demonstrate a lack of understanding of this kind by attempting to get into a car that is just about the size of their foot. Of course, this is an obviously impossible task and the toddler will not be able to fit their body into the small car no matter how many times he or she tries. This kind of error is a scale error. DeLoache *et al* explain that scale errors demonstrate a clear dissociation between the perception of the body and motor behaviour[96]. In more precise terms, the infants' failure to apprehend the significance of their body size relative to the car feeds directly into their failed attempt to make an appropriate motor choice. The interesting point is that personal level appreciation of the body in terms of its scale and motor potentiality is important for determining whether an infant is capable, or not, of making appropriate motor choices. As such, it is crucially important for shaping and constraining the motor repertoire of the developing infant. Across

---

[96] See DeLoache *et al,* 2004 for seminal work on scale errors.

the course of development there is ample opportunity for the developing infants' experience of its body to continually readapt to the changing motor capabilities of its body, and vice versa. For example, the young baby at only a few months old who can only lie on his or her back will have a much different visual experience of their legs to the young toddler learning to walk for the first time. The former infant will typically be able to see her legs by raising them up towards their torso, while the letter infant will typically look down to look at her legs as one moves in front of the other. Likewise, the visual experience an infant has of their hand meeting a cup when they can reach for and grasp it at around 8-months-of-age will be different to the experience a younger infant who is yet to develop the ability to extend their hand towards objects in an intentionally purposive way.[97]

The upshot is that within recent empirical work that's currently going on in developmental psychology, the body image and schema distinction has the potential to make an important contribution to a story that can help explain the underlying processes which underpin young infants' apprehension of the interdependent relationship between the configuration of its body and its motor capabilities has an important role to play in the development of its perceptual apprehension of its body and motor repertoire, and how incremental changes to their interrelations drive appropriate changes to the perception and movement of their body. Telling this kind of story requires a theoretical framework that will allows us to explain the distinct features of the the body image and the body schema which can then be brought together to explain their interrelations. The natural question is what theoretical framework developmental psychologists should appeal to. The proposal I offer in the remaining sections of this chapter brings together the neuroconstructivist model of development that was discussed in chapter one and Shaun Gallagher's analysis of the developmental relationship between the body image and the body schema.

6.4. The Neuroconstructivist Approach in Developmental Psychology.

---

[97] Bremner & Spence, 2012.

The rationale of the neuroconstructivist model of development was outlined and explored in brief detail in chapter one of this thesis with respect to the neuroconstructivist interpretation Holmes & Spence support in relation to their developmental studies on the stages at which young infants lack, and subsequently come to possess, the ability to visually track the spatial rearrangement of their body parts by way of a process that involves, in part, Bayesian style predictions. It is now time to consider the aims of neuroconstructivism in some more detail. As the name suggests, neuroconstructivism[98] is, in part, based on the constructivist outlook from Piaget (Piaget, 1955) which seeks to understand the transformative effects of experience-dependent knowledge on the gradual acquisition and progression of the mental and cognitive skills of developing children that are acquired through their active engagement with their (physical and social) environment. The 'neuro' in neuroconstructivism refers to its commitment to the idea that the structural configuration; functional segregation; integration and specialisation of specific regions in the brain are built up out of a process by which the capabilities of individual neurons and classes of neurons is determined by the interactions between neighbouring neurons. In turn, the mental and cognitive capabilities that these neural populations support is further shaped by the embodied entity in which they are enveloped (i.e., what kind of body the agent has) and the interactions of the body in the physical and social environment. Ultimately, neuroconstructivists endorse a collective approach to understanding mental and cognitive approach that draws upon the insights from various disciplines including genetics, neuroscience, biology, psychology, embodied and social cognition and phenomenology. Neuroconstructivism is a rich and multifaceted framework, but for present purposes it is important to note that an important and distinctive hallmark of neuroconstructivism is that it places strong emphasis on the need for a causal explanation of the driving forces which are responsible for the unfolding of mental and cognitive phenomena along developmental timelines. Thus, we are told:

> The biggest challenge facing developmental psychologists is *to link these observations into a developmental trajectory* and *to explain the causes of developmental change* that allow the child to progress from one set of abilities to another, more complex one.
> (Mareschal *et al*., 2007, p.389 my emphasis).

---

[98] See, e.g., Quartz & Sejnowski, 1997, Karmiloff-Smith, 1998 and Mareschal, Johnson Sirois, Spratling, Thomas & Westermann, 2007.

The importance of this requirement stems from a dissatisfaction with what neuroconstructivists call the standard 'snap-shot' approach in developmental psychology. As the label suggests, this involves psychologists taking a very brief 'snap-shot' of the capabilities, or lack thereof, of children at different stages of development for comparison. For instance, Piaget demonstrated that children at 7-, but not at 10-years-of-age, lack the concept of conservation which is important for understanding that the objective configuration of objects remain the same even if their perspectival properties change. Youngsters of 7-years-of-age lack the understanding that even though the round plate looks elliptical, it remains round. However, by their tenth-year children have acquired this perceptual understanding (Piaget, 1955). The claim of the neuroconstructivist is that this approach is informative in one sense, but uninformative in another. It is informative because we at least know the ages at which there is a difference in the child's understanding of the relationship between what Locke would call the primary and secondary qualities of objects. The primary qualities are those attributes of the object which remain constant no matter the conditions under which the object is perceived, i.e., the actual shape and size of the object. Therefore, primary qualities are perceiver independent. Secondary qualities are perceiver-dependent, e.g., the (perceived) colour; shape; taste of the object, etc. (Locke, 1689). However, the information such studies provide is also not very informative in the sense that it doesn't account for *how* the relevant knowledge is acquired for it only tells us that this acquisition occurs at some *as-yet-unspecified* time between 7- and 10-years-of-age. A three-year age gap is a big leap in terms of development, and in the absence of an understanding of the processes that drive this change the 'snap-shot' approach doesn't tell us very much. As such, neuroconstructivists argue that developmental psychology is currently subject to an important deficit that can only be rectified by:

> […] *causal theories regarding what makes complex behaviours emerge* These theories need *to explain behaviours on multiple time scales*. They must explain *how and why behaviours unfold as we observe them in real time*, *as well as how they unfold in developmental time*. To do this, we need more than just a very detailed description of the behaviours that can be observed at any point (Mareschal, *et al.* 2008, p. 331, my emphasis).

Thus, to tell a sufficient story about the relationship between the body schema and the body image we need to explain the mutual transformative effects that each has on the other along developmental timelines and in real time. To do so, I shall revisit Shaun Gallagher's

behavioural analysis of the relationship between the body schema and the body percept during motor learning and everyday action.

6.5. Gallagher's Developmental Behavioural Analysis of the Body Percept and the Body Schema: A Brief Reminder.

In chapter two of the thesis it was observed that in addition to showing how dissociations between the body image and the body schema can be made, Shaun Gallagher also provides an analysis of the reciprocal relationship between them during motor learning. Let's return to that analysis. During motor learning developing infants and mature adults build their motor repertoire by adding a novel motor skill (or set of motor skills) to their pre-established set of motor skills. Through the process of learning and developing new motor skills the developing infant faces the additional challenge of understanding how the continually changing proportions of its body directly affect its motor capabilities and constraints. The important issue that emerges out of Gallagher's line of thought is why learners exhibit an increased reliance on their body percept during the development of his or her motor skills. I think we can make better sense of this observation in terms of free-energy minimisation and active inference.

To tell a story about the developmental significance of the body percept for the development and transformation of the motor behaviours of the body schema is to tell a more general story about the processes through which motor processes are developed and transformed by the influence of personal-level perceptual processes. If the prediction-error minimisation model is correct, then it is through the process of establishing an equilibrium between the perceptual predictions about our body and sensory feedback that gives us meaningful perceptual access to the state and significance of our body.

Under Gallagher's interpretation, the body percept is underpinned by the sensory feedback that's generated by the body, e.g., visual, proprioceptive, tactile and auditory information about the perceptible features of the body. In terms of the prediction error minimisation framework the body percept is a distinct set of sensory and sensorimotor predictions that pertain exclusively to the perceptible features and motor activities of the body. For ease of exposition, let's agree that the body has a distinct *cross-modal predictive profile* which is

made up of the class of predictions which are exclusively about long-term properties of the body (e.g., its canonical configuration), its short-term properties (e.g., its current postural position) and any other relevant perceptible properties (e.g., motor generated feedback.) As such, the body percept is a class of hierarchically structured predictions that separate the body from other objects in the environment that are characterised by other distinct cross-modal predictive profiles.) The job of the body percept is to provide us with meaningful perceptual contact with our body, such as what it looks like in general; what it looks like when it moves; what it is currently like in its current postural positions; how it differs from any other object in space, and so forth. Thus, to tell a story about the body percept is to tell a story about how multisensory processes establish, transform and maintain perceptual awareness and recognition of the body. In turn, making a distinction between the body percept and the body schema requires an account of the differences between them in terms of the distinct predictions they elicit.

Under active inference successfully generating novel actions requires a lot of fine-tuning for the predictive capacities of the recognition density to come into line with the trajectory of sensory inputs that are generated by previously untested movements. In such cases, the predictive capacities of the recognition density are under-developed such that it can't reliably predict the sensory consequences of the movement(s) with any degree of accuracy. During ordinary motor learning, there is an extensive period of fine-tuning between the predictions elicited by the recognition density and sensory feedback that's produced by the movement. When the predictions and sensory inputs are brought into equilibrium (i.e., the model can accurately predict the variable sensory consequences of the movement) they can be performed below threshold of consciousness. This applies just as much to the growing infant who is learning to crawl and walk as it does to the mature adult learning to dance. In cases where the individual is carrying out an action that could result in possible injury or death, the potential for producing a sensory prediction error remains high until the action has been successfully carried out. This could perhaps explain why the conscious deliberation and monitoring of movement throughout the duration of the movement. The key point is that under active inference learned behaviours are tested and mastered in accordance with personal-level expectations until the repeatable sensory feedback they generate and the predictions elicited by the recognition density are brought into equilibrium. In cases where one is entering a high-risk situation, it helps to proceed at a slower rate than usual and consciously keep track of our movement for it isn't always easy to predict the consequences of action when we find ourselves in risky situations. In this case, it makes sense to explain the

behaviour in terms of the desire to avoid injury as opposed to the fulfilment of certain expectations about the sensory consequences of action. From a free-energy perspective the more a behaviour minimises unpredictable (i.e., erroneous) effects that are not accounted for by the generative models the higher the prior likelihood that it will minimise free injury and, as we know, minimising free energy equates to securing the optimal existence of the agent in its environment.

  This important point proceeds from the common-sense observation that learned behaviours start off unsteady and imprecise but, through practice, become increasingly more accurate, refined and effective. Once mastered, motor tasks can be performed without conscious deliberation and planning. As we have observed before, a mature adult can walk, reach for objects and pick them up, walk upstairs etc., without having to consciously think about how to do the bodily movements that make up those kinds of actions. By this stage, the behaviours have already been deferred to the body schema to be produced and regulated below the threshold of personal-level consciousness (Gallagher, 2005, pp.45-52). However, the fact that mature humans perform these tasks with relative ease and little conscious reflection makes it easy to underestimate the protracted and extensive sensorimotor learning period that preceded this mature state. If able, the reader will most likely be able to recall the many times they ended up with grazed knees and elbows that go along with learning to ride a bike and, except for beginner's luck, the many failed attempts that took place prior to being able to throw a netball into the net with a sufficient frequency of success. As Gallagher points out, when we are learning to do these activities successfully, we tend to pay more close attention to our body, such as the overall posture of our body when we are sat on the bike. In this case the trunk of our body must stay fully erect and tense for if our body sways too much to one side we will likely fall over. The same applies to the speed at which the hand is extended towards the net as the ball is released towards it. If one releases the ball too forcefully when one is near the net, the ball may just recoil off the backboard. However, if one is much further away more exertion and force is required to propel the ball from a greater distance. The point is that effective motor learning requires more personal-level guidance relative to motor behaviours we have already mastered. The same applies to the more fundamental universal motor skills that underpin advanced motor skills such as the ability to sit and stand up straight, walk, jump and so on. Likewise, developing infants tend to pay close attention to movement of their bodies during the development of their motor skills. Think of the infant learning to walk for the first time who pays close attention to the movement of her feet as she while being held up by her caregiver.

6.6. The Developing Body Image: An Analysis of Canonical and Postural Remapping

The preceding line of thought puts is in a better position to analyse the developmental studies from Holmes & Spence in their *2012* apply to their 2008 developmental studies to explain how the body percept might be established through knowledge driven perceptual inference. Recall that the studies employ the preferential looking paradigm to investigate whether the visual responses of young infants (i.e., eye tracking) can accurately keep track of tactile sensations on the body parts as they move across their visual space[99]. The upshot of such studies is that infants at six-and-a-half months of age do not exhibit the ability to visually track the movement of their body whereas infants at the age of ten-and-a-half months of age can do this.

The underlying gist of the story is that, as mature adults, we understand that when our body parts are in their canonical (i.e., typical) spatial positions they occupy the respective spatial region of our visual field. The left side of our body occupies the left region of visual space, and vice versa for the right side of the body and visual space. However, if the arms are crossed over to the other side such that, say, the right hand now occupies the left side of visual space, any activity of, or in relation to, the right hand is still registered in terms of the activity of the right hand or in relation to it. Therefore, the movement of this hand is still registered as movement of the right hand, not the left. Likewise, if a fly were on this hand the tactile stimulus it generated will still be registered as coming from the right hand. Therefore, when my right hand moves it generates visual stimuli that occupied the left region of visual field. As such, I may glance to my left instead of my right if, say, I want to work out what is causing the tickly sensation on my right hand as a fly walks across the surface of my skin. Holmes and Spence established that the ability to readapt visual responses in line with the movement of the body parts is present in infants at ten-months-of-age, but not at six-and-a-half months. The six-month-old infant will glance to their right because this is the typical way to gain visual access to the right hand as they fail to appreciate that the movement of the hand requires a different visual response (i.e., looking left).

---

[99] For a more detailed reminder of the details of the study refer to section 3.6 of chapter 3 in the thesis (pp.61-65).

Recall that from a general Bayesian perspective the visual behaviour of the ten-month-old infant and the mature adult is underpinned by the antecedent expectation that looking to the left when my right hand is positioned at the left side of my body will give me perceptual access to my right hand. Under PEM the story would, initially, follow along the same lines. It would be agreed that the crossover of the right hand from the right side of the body to the left side provides updated proprioceptive information about the position of the right hand. This causes a change in the antecedent visual-proprioceptive expectations about how perceptual access to the right hand can be brought about. Therefore, the cause of the change in conditional expectations about the sensory consequences of events in relation to the right hand is the switch to a more complex set of visuo-tactile-proprioceptive predictions. Additionally, the appropriate changes in motor behaviour, i.e., glancing to the left instead of the right can be explained in terms of active inference. The stimuli produced by the fly landing on my hand would instigate a generative model which would predict the cause of the stimulation. In turn, the perceptual system would engage the oculomotor motor system to direct the eye in the direction of the tactile stimuli on the body to sample the sensory data to verify that the predictive power of the generative model is accurate (i.e., that it is a fly on the hand) or detect any relevant prediction errors that need to be suppressed by further perceptual or active inference. Looking to the right in this case will produce a prediction error because the hand won't be visually located. In turn, the incoming prediction error should generate alternative motor behaviours until the prediction error is quashed by looking to the left.

Recall that it is precisely because this adult-like behavioural response is present in ten-month-old infants, but not those of six-and-a-half-months, Holmes and Spence posit the existence of two distinct multimodal 'body representations'. The 'canonical' body map represents the canonical configuration of the body and is used to localise sensations on the body when they are positioned in their standard positions. Then there is 'postural remapping' which localises sensations on the body in accordance with the changing positions of the body parts. They speculate that the latter emerges out of the former at some point between six-and-a-half months and ten-months- of- age. Under PEM the canonical body map and postural are two distinct but interrelated perceptual generative models which encode probabilistic predictions about the location of sensations on the body part, and the establishment of the latter is driven by active inference. For example, it's anyone's guess as to precisely when the infant develops this understanding between the specified age ranges that Holmes & Spence investigate. However, when the six-and-a-half-month old infant looks to the right under the expectation that they will see their right hand only to see that their hand is not there this will

generate a prediction error which in turn instigates further active inference, such as more looking behaviours until the right hand can be seen at the left side of their visual field. In turn this engenders the establishment of an alternative generative model which elicits more accurate visual-proprioceptive predictions that underpin more appropriate behavioural response next time around.

6.7. Understanding Scale Errors in Terms of Active Inference.

Under active inference, scale errors provide an important opportunity to provide relevant sensory and sensorimotor prediction error(s) that are important for updating the developing infant's generative model with a new prior (e.g., the prior likelihood of the body fitting into enterable objects of the same size of the car is nil). This should inhibit the behaviour for under active inference the sensory and sensorimotor prediction errors may instigate further attempts to enter the car. However, each failed attempt cannot satisfy the initial predictions that are responsible for eliciting this behaviour because the body will never fit into the scaled down car. The overarching point is that it is through interactive experience the developing infant develops the predictions and priors necessary for enabling a suitable array of behaviours that are sensitive to the constraints placed by its body and environment.

6.8. Gallagher's Behavioural Analysis of the Body Image and Schema in Real Time.

Gallagher tells us that there are two circumstances under which the movement of the body requires conscious perceptual guidance that is provided by the body percept. In such cases the body percept assists the body schema to ensure that the body moves successfully through the environment. We have looked at the first case of motor development. The second pertains to conditions in which an individual must manoeuvre his or her body in situations that are potentially risky and/or fatal. To begin let's recall the proposition that it is the job of the body schema to keep the body poised for swift and precise fine-grained interactive motor

behaviours, such as avoidant responses like ducking the head to avoid a collision with a low hanging tree branch and quickly raising the hand to swat wasps away to avoid a painful sting. However, there are occasions where it seems that the sub-personal body schema is not enough to initiate behavioural responses that get an individual out of an unwanted and potentially hazardous situation. Gallagher tells us that when an individual is walking along an icy pathway or a high ledge he or she must proceed with caution, lest they make a serious miscalculation that could result in a serious fall which could leave them seriously injured or dead. To do so, it makes sense to switch from the 'auto-pilot' mode of the body schema to consciously guide our actions using, in part, the sensory processes of the body percept. In such a situation, we carefully slow down our pace, deliberate each step before we carry it out and visually monitor our progression in a similar way to the unsteady and accident prone infant who is learning to walk for the first time who closely watches her legs as she makes each step or the mature adult learning to dance for the first time (I'll come back to this point).

This immediately raises the issue of why we defer to our body percept in such situations when, as we have just observed, often the movements produced by the body percept are less effective than those produced by the body schema. To a first approximation this may seem like a counterintuitive strategy but it is not. First, the motor procedures produced and regulated by the sub-personal processes of the body schema can only be considered optimal when they produce quick, precise, and efficient motor responses that don't make as much demands on our cognitive workload and can reliably eliminate any impending threats to the same degree as the conscious guidance of actions. In some cases, this might be true. For instance, swatting away insects with our hand is usually enough to deter them. However, slowing down the rate at which we move could bring about an adaptive payoff in situations where the stakes are high. Moving too swiftly might cause a motor error that requires a corrective measure (or sequence of corrective measures) to rectify it. Such compensatory measures will make more demands on our cognitive resources. In turn, the shock of falling or injuring oneself will place further demands on our cognitive apparatus. For instance, the adrenaline rush and come down period that comes with the shock of falling, the injuries sustained and the metabolic resources required to get back up. Therefore, it is surely worth the cognitive effort required to calculate the next manoeuvre at each stage in high-risk situations such as these to avoid the consequences and after-effects of motor mistakes.

This provides two important insights when it comes to thinking about the importance of the body percept and body schema for motor performance. The first is that having a personal-level body percept, on its own, is insufficient to optimise motor performance. Having a body

schema that can function without the need for ongoing personal-level deliberation and monitoring optimises behaviours because the behaviours it produces are quicker, more precise[100] and place less demands on our cognitive resources. Nevertheless, this doesn't mean that the body schema is sufficient, on its own, to effectively manoeuvre the body in space; rather, having a body percept that can take over the role of the body schema to guide the movement of the body is adaptively advantageous, for the reasons specified above. The fundamental question is why we alternate between the body image and the body schema in this way. One such case was the deafferentated patient Ian Waterman. After a sensory neuropathy Ian lost all control over his motor behaviours. He lost the sensation of touch (i.e., sensation of pressure on the skin) and proprioception (i.e., the sensation of the position of his limbs). In short, Ian suffered a complete loss of his body schema and a partial loss of his body percept for he still his access to his visual percept. In other words, he still receives the visual information about the position and movement of his body which structures his visual awareness of his body. After extensive rehabilitation Ian, recovered his motor skills. However, now he is completely reliant on his visual percept. Following the considerations that were discussed in the previous chapter it seems that having a body percept that can perform the same functions as the body schema is adaptively beneficial for Ian in the sense that without his visual percept he wouldn't have been able to regain his motor skills at all. In other words, not only does this case suggest that there are dissociations to be made between the body image and the body schema, it also suggests that having two functionally equivalent systems which can take over if the other fails. However, Ian's behaviours are far from optimal. His movements are slower, less precise, and make more demands on his personal-level resources (through the constant requirement of visually fixating on his body.) The natural question is why this is the case. Ordinarily, you or I can alternate between personal-level visual guidance and sub-personal motor planning, so why can't Ian? It is difficult to give a clear-cut analysis, but here are my speculations from the perspective of the prediction error minimisation approach.

Ian has made the switch from a malfunctioning system to a degraded system as he can only rely on visual proprioceptive awareness. To re-establish his motor skills Ian has had to form entirely novel ways of performing the same action. What's more is that only one sensory modality is providing the information necessary to move his body, i.e., visual information.

---

[100] Remember the studies from Jeannerod & Jacob (2003) which showed that even when an object cannot be seen, the way that the body schema instructs the formation of the hand into a grasp is sensitive to the objective three-dimensional configuration of the object.

The studies from Friston show that ordinarily the training of one modality is sharpened up by its coupling with other modalities during different tasks and under different conditions. In other words, the predictions a single modality carries are fine-tuned to be made more precise to compensate for its context-dependent and task-specific limitations and so that it can perform the job of the others in their absence. On top of this, in chapter 3 we observed that perceptual estimates are inherently probabilistic in virtue of the inherently noisy nature of signalling. As such, Ian is relying on a single source of information which doesn't offer the same robustness as two or more cross-modal signals engaging in active inference would when they are in the business of working together to quash modality-specific and cross-modal prediction errors. Therefore, Ian is missing two crucial components of the normal motor learning process: cross-modal prediction and exposure to cross-modal prediction errors.

The second important point is that the role of the body percept in conditions which are likely to pose a serious risk of injury can be explained in terms of free-energy minimisation and active inference. Behaviours for which the outcomes are less predictable and less certain present the real risk of increasing free-energy. For example, when we're walking on ice each time one foot is raised off the ground we are at risk of slipping, and the same applies to each time our foot comes back into contact with the icy surface as we place one foot in front of the other. The same applies to walking along a high ledge. All it takes is a miscalculation about how one should proceed (e.g., placing the foot on unstable terrain) for one to be seriously injured and/or dead. As such, it is harder to maintain equilibrium between the cross-modal predictions that are driving the behaviours and the consequences those behaviours produce. The common thread that runs through both conditions under which we may recruit our body percept is that they carry a greater prior likelihood of generating a big prediction error, by which I mean they raise the possibility of the brain making a mistake in its predictions about the sensory consequences of the movement that could have major consequences for the individual. In turn, big prediction errors could potentially compromise the optimal function of the agent or ultimately cause its death. Furthermore, the prior likelihood of various kinds of prediction errors coming in at a much faster rate is also higher in situations where the brain is less able to calculate the outcomes of behaviours to a reasonable degree. Under the free-energy principle the actions of the agent are produced to test out the predictive power of its generative models whilst respecting its morphological constraints and homeostatic requirements. Choosing any potential course of action that could result in serious injury or death violates the free-energy principle.

This is an admittedly rough-and-ready sketch of how Gallagher's analysis can be illuminated when it is reconsidered in terms of active inference and free energy minimisations, but it provides us with just enough to provide a plausible story about the neural underpinnings which drive the switch from the body schema to the body image. Once it is built up to a sufficient degree the body schema can successfully produce and regulate the context-and task-specific motor performances of the body without the need for personal-level deliberation and/or monitoring because each time it initiates a motor behaviours provides an opportunity to test their free-energy minimising potential. Doing this at a sub-personal level is the optimal thing to do as it means that the agent is not dependent upon their more cognitively demanding personal-level capacities to regulate their behaviours. If this is indeed the case, it makes adaptive sense to *only* elevate to the personal level to assess the situation of the body in unpredictable and/or hazardous conditions to calculate the nature and severity level of the risk and select the movements that are most likely to successfully navigate the body in such conditions, i.e., by minimising and/or removing the risk of imminent threats. Thus, the difference between the threat of overhanging tree branches and icy pathways is that the former presents a threat that can be dealt with because the body schema is equipped with a proven successful behavioural strategy to deal with it (e.g., ducking the head.) Thus, it is expected (i.e., predicted) that that behavioural strategy will remove the threat of a collision. The consequences of walking on an icy pathway are less predictable, and more dangerous, and the behavioural strategies the body is equipped with do not guarantee that the individual won't fall with the same level of expected certainty. In short, whether the body schema or body percept is deployed to guide the movement of the body depends on how predictable the consequences of the behavioural strategies produced are going to be.

6.9. An Interpretation of the Rubber Hand Illusion from the Perspective of Predictive Processing.

At various stages of this thesis we have observed the invaluable insights that can be gleaned from an examination of the effects that can be brought about by altering cross-modal couplings in various ways. This strategy gives us some leverage to work with when it comes to examining and explaining how multisensory processes optimise the perceptual and motor processes they support. The first instance of this kind of methodological strategy was outlined in chapter one of the thesis in relation to the so-called 'rubber-hand illusion' where

participants report experiencing sensations on a rubber hand. As such, the rubber-hand illusion is important for analysing the way in which cross-modal sensory relations shape our perceptual apprehension of the boundaries of our body. Recall that the experiment involves a participant who has one hand occluded behind a screen and a realistic rubber hand is substituted in its usual place. At the same time, their occluded hand and the rubber hand receive identical tactile stimulation. When asked to locate the sensation on their body part most participants report that they feel the location of the sensation on the rubber hand, not their occluded hand.

I mentioned along the way (in chapters one and three) that a common explanation for the fact that visual input overrides tactile-proprioceptive input in this case is because vision remains the dominant modality when it comes to localising the position of our body parts, and the location of sensations thereon. This assumption squares nicely with Ernst's maximum likelihood principle from chapter three for we have good reason to believe that whenever redundant visual and tactile inputs are coupled to provide a cross-modal estimate of the size of an object, the visual signal will always be given a higher weighting than the haptic signal. In this respect the rubber hand illusion provides a nice illustration of how our perceptual understanding of the boundaries of our body is shaped and altered by cross-modal relations about the location of sensations on the body. But we are still far away from an explanation as to why the visual modality overrides tactile inputs and perhaps this is best explained in terms of perceptual inference as it is construed under PEM. When the occluded and rubber hand receive tactile stimulation, this creates a conflict between the visual stimuli and tactile stimuli about the location of the sensation (i.e., on the rubber vs. occluded hand, respectively) and proprioception (about the location of the occluded hand). The experimental design that produces the rubber hand illusion thus creates a conflict between visual-tactile inputs about the location of the stimuli and tactile-proprioceptive inputs about the location of the stimulus of the real hand. To quash the prediction error, the visual-tactile predictions about the location of the stimulus on the rubber hand overrides the tactile- proprioceptive predictions because of the prior likelihood that the visual signal is the more reliable signal about the location of the stimulus, even though this generates a perceptual error (i.e., the impossible judgement that the subject feels the sensation on the rubber hand). However, perhaps this is the small trade-off we pay for the fact that multisensory processes make perceptual and motor processes more robust than unisensory processes in general.

Conclusion

To be able to tell a story about the body image and body schema distinction in accordance with the requirements set by neuroconstructivism one must be able to provide an explanation of the nature of the driving forces, i.e., causal, behind the perceptual and motor capabilities they support in a way that will help us understand how and why they unfold in real time, and over the course of development. In turn, to tell a story about the respective processes which characterise the body image and the body schema, as well as the relationship between them, is to tell a story about the respective nature of perceptual and motor processes and the unifying principle(s) which unite them to support body perception and body movement. In chapter four such a story was provided under the Bayesian inspired hierarchical predictive processing models of perceptual and motor processing. Perception, under this framework, is a process of knowledge driven inference through which an individual gains meaningful perceptual access to their environment. More precisely, pre-acquired knowledge about the causal structure of the body and environment is encoded in hierarchically structured generative models) which elicit probabilistic predictions about the cause(s) of incoming stimuli. In the case of action, motor behaviours are produced by hierarchical generative models which elicit probabilistic sensorimotor predictions about the trajectory of sensory data that should ensue at each stage of the movement trajectory. Therefore, insofar as perceptual and motor processes are distinct it is because they deal in different predictions. However, they are united a common goal: prediction error minimisation. Under the active inference principle, the perceptual and motor processes of the agent are united to optimise the predictive capabilities of the generative models by actively sample its sensory environment to quash prediction errors. Remember that an agent whose predictions can best account for all the sensory data its activities induce is at equilibrium with its environment and, under the free-energy principle the optimal existence of classes of animals, and their individual members, necessitates that their brains encode a generative model (or sets of models) about the predictable (and random) effects that its activities have on its internal economy at all levels of functioning. This gives us two different timelines from which to explain how perceptual and motor processes develop and transform. By applying this framework to Gallagher's behavioural analysis of the body schema and body percept distinction I have illustrated the ways in which we can provide a neuroconstructivist picture of how the

behavioural transformations which are driven by changes to the perceptual apprehension of the body over two different timelines (i.e., developmental and in real time).

Thesis Conclusion

The driving motivation of this thesis has been to explain the nature of the optimisation strategies that determine the contribution multisensory processes make to whatever mental and cognitive phenomena they support. I used the pre-existing discussions in cognitive psychology and philosophy as a platform from which to illuminate the key questions which shape contemporary multisensory integration research in cognitive & developmental psychology and philosophy. Across the first two chapters we saw that the initial strategy is always the same: to clear up the conceptual confusions pertaining to the body schema. Of course, the methodological approach is quite different for whilst psychologists attempt to pin down the neurological underpinnings of the body schema, philosophers seek to provide a clear conceptual analysis of this concept. Either way, the first two chapters of this thesis illustrate that in each camp confusions over the body schema remain, objectionable interpretations and assumptions remain unchallenged, and the value of the distinction remains unacknowledged. By using Shaun Gallagher's analysis I hope to have taken reasonable steps towards remedying these issues by continuing his project of clearing up the conceptual confusion; challenging the unwarranted presuppositions of his opponents about the nature of the body schema and the future of the distinction between the body image/schema in cognitive science, and showing how the distinction is a useful explanatory tool for working towards a neuroconstructivist model of the developmental relationship between the body image and body schema. This story is by no means complete, but it is, I believe, the beginnings of a more illuminating story about how multisensory integration shapes and transforms body perception, motor control and their interrelations.

References

Alais, D & Burr, D (2004) The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*. 14(3) pp. 257-262.

Banks, M & Ernst, E (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*. 41(5), pp. 429-33.

Bernardo, A, J.Carlin, D. Dunson, S.Stern, Vehtari.A (2004) Bayesian data analysis, second edition, Chapman & Hall: CRC Press.

Berkeley, G (1709) An essay towards a new theory of vision.

Bremner, A, Marivata, A & Spence, C (2003). Multisensory integration and the body schema: Close to hand and within reach. *Current Biology*. 13, pp.531-539.

Bremner A, Lewkowicz, J & Spence, C (2012). Multisensory Development. Oxford: OUP.

Botnivik, M (1998) Rubber hands feel touch that hand see. *Nature*. (19), pp.669-756.

Calvert, G, Spence, C & Stein, B (2004). The Handbook of Multisensory Processes. Cambridge, MA: MIT Press.

Carruthers, G (2008). Two types of body representation and the sense of embodiment in *Consciousness & Cognition,* 17, pp. 1302-1316.

Carruthers, G (2009). Is the body schema sufficient for the sense of embodiment? An alternative to de Vignemont's model in *Philosophical Psychology,* 22(2), pp. 123-142.

Clark, A (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioural and Brain Sciences*, 36(3). pp. 181-204.

Clark, A (2014). Perceiving as Predicting. In M. Mohan, S.Biggs, & D.Stokes (Eds). *Perception and its modalities*. New York: Oxford University Press.

Clark, A (2016). Surfing Uncertainty: prediction, action and the embodied mind. New York: Oxford University Press.

Clark, A & Pickering, M (2014). Getting ahead: forward models and their place in the cognitive architecture. *Trends in Cognitive Science* 18(9): pp.451-456.

DeLoache, J (1998). Scale Errors offer Evidence for a Perception Action Dissociation in Early Life. *Science.* Vol. 304. Pp. 1027-1300.

Descartes, R (1641). Meditations on First Philosophy reprinted in Cottingham, J (ed) for Cambridge texts in the history of philosophy (2006).

De Vignemont, F (2006) A review of Shaun Gallagher's how the body shapes the mind. *Psyche* 12(1).

De Vignemont, F (2010) The body image and schema distinction: pros and cons. *Neuropsychologia,* 48(3), pp.669-80.

Ernst, M (2006). A Bayesian view of multimodal cue integration. In *Perception.* 131(6), pp.105-131.

Friston, K (2009). The free-energy principle: A rough guide to the brain. *Trends in cognitive science,* 13, pp. 293-301.

Friston, K (2010). The free energy principle: A unified brain theory? *Nat. Rev. Neurosci,* 11(2), pp. 127-138.

Friston, K (2011). What is optimal about motor control? *Neuron,* 72, pp. 488-498.

Friston, K (2011). Embodied inference: or 'I think therefore I am, if I am what I think;. in W.Tschacher, & Bergoni, C (Eds) *The implications of embodiment, cognition and communication* (pp.89-125). Exeter, Uk. Imprint Academic.

Fodor, J.A (1983) Modularity of Mind, Cambridge, MA: MIT Press

Gallagher, S (1986) Body image and body schema: a conceptual clarification. *Mind & Behaviour* 7(4), pp.541 – 554.

Gallagher, S & Coles, J (1998), Body image and body schema in a deafferentated patient. *Mind & Behaviour* 16(4), pp. 369-390

Gallagher, S (2005). How the body shapes the Mind. New York: Oxford University Press.

Gibson, J.J (1968). What gives rise to the perception of motion? *Psychological review,* vol.75, No.4, 335-346

Gibson, J.J (1966). The senses considered as perceptual systems. Boston: Houghton Mifflin.

Graziano, MSA, Yap, GS, Cross, CG (1995) Coding of visual space by pre-motor neurons. *Science.* (266), pp. 1054-157.

Graziano,MSA, Reiss, LA, Cross, GC (1999) Complex movements evoked by microstimulation of precentral cortex. *Neuron.* (34), pp. 851-861.

Graziano, MSA, Hu, XT, Gross, CG (2002). Visuospatial properties of ventral premotor cortex. *Journal of Neurophysiology.* (77), pp. 2268-92.

Grigor, P (1998). Aristotle on the common sense. New York: OUP.

Green, DM, Swetz JA (1988). Signal detection theory and psychophysics. Pensinsula Publishing.

Helmholtz, H (1978). The Facts of Perception. Wesleyan University Press.

Holmes, N & Spence, C (2006). Beyond the body schema in *Human body perception from the inside out*, Knoblich, G, Thornton, I, Grosjean, M & Shiffrar, M (Eds). Oxford: OUP.

Holmes, G & Head, H (1911-12), Sensory disturbances from cerebral lesions. Brain *34.* P.187.

Holmes, G & Head, H (1920). Studies in neurology. London: Oxford University Press.

Hohwy, J (2007) Functional integration and the mind. *Synthese,* 159(3): pp.315-328.

Hohwy, J (2012) The predictive mind. Oxford University Press.

Hurley, S (1998) Consciousness in action. Harvard University Press.

Iriki, A, Tanaka M, Iwamura Y (1996). Coding of modified body schema during tool use by macaque postcentral neurons. *Neuroreport*, 1996a, pp.2325-2330.

Jacob, P & Jeannerod, M (2003). Ways of seeing: the scope and limits of visual cognition. Oxford: Oxford University Press.

Jeannerod, J (2007). Simulated actions share first and third person representations. *Brain Research* 1130(1): pp.125-29.

Karmiloff-Smith, A (1998). Development is key to understanding developmental disorders. *Trends in Cognitive Science* 2(10): pp.389-98.

Karmiloff-Smith A, Thomas M, Westermann, G (2007). Neuroconstructivism. In *The Handbook of Cognitive Development*. Oxford: Blackwell.

Lhermitte, J., & Tchehrazi, E. (1937). L'image du moi corporel et ses déformations pathologiques. L'encéphale, 32, 1–24.

Locke, J (1698). An essay concerning human understanding. T.Tegg & Son.

Mandik, P (2005). Action-oriented representation in Brook, Andrew & Annens, K (eds.) *Cognition and the Brain: the philosophy and neuroscience of movement.* Cambridge: Cambridge University Press.

Mendoza, J (2011). The Handbook of Neuropsychology, vol.2, 8[th] edition, Elsevier.

Meredith, A & Stein, E (1983). Interactions among the merging senses in the superior colliculus. *Science,* 221, 389-391.

Meredith A, Stein, E, Wallace, M (1998), 80(2). Multisensory integration in the superior colliculus in the alert cat. *Journal of Neurophysiology.* 80(2), pp.1060-1010.

Meredith, A & Stein, E (1993). The merging of the senses. MIT press.

Millikan, R (1983). Language, thought and other biological categories: new foundations for realism. Nous 21(3): pp.430-434.

Noë, A (2005). Action in Perception. MIT press.

Noë A, O'Regan K. (2001). A sensorimotor account of vision and visual consciousness.

Paillard, J (1999). Body image and body schema: A double dissociation. In *Motor Control Today and Tomorrow,* Gantchev EN, Mori S, Massion J (Eds).

Reid, V (2014). How infants detect information in biological motion in Brownell, A & Slaughter, V (eds): *Early Development of Body Representations*. Cambridge Studies in cognitive and perceptual development. Cambridge University Press.

Reinoso, R, Carvalho, F (2016). The influence of soundscapes on the perception and evaluation of beers. *Food Quality and Preference* (52), pp.32-41.

Rosetti, Y, G.Rode & Boisson, D (2001). Numbsense: A Case Study and Implications in B. De Gelder, E. De Haan & C.Heywood (Eds) *Out of Mind: Varieties of Unconscious Processing.* Oxford: Oxford University Press.

Slaughter V, Brownell, A (2012). Early Development of Body Representations. Cambridge Studies in cognitive and perceptual development. Cambridge University Press.

Spence, C (2014). Noise and its impact on the perception of food and drink. *Flavour*, pp. 3-9.

Spence, C (2015). On the psychological impact of food colour. *Flavour*, pp. 4-21.

Stein, E & Stanford, J (2008). Multisensory Integration: current issues from the perspective of the single neuron. Nat Rev Neurosci, 9(4), pp. 255-266.

van Beers, R.J Sittig AC, Denier van der Gon (1999). The precision of proprioceptive touch sense. *Experimental Brain research.* 122, pp. 367-377.

Tsakiris, M (2014). The embodied mini-me: tracing the development of body representations and their role for self-awareness in Brownell, A & Slaughter, V (eds): *Early Development of Body Representations.* Cambridge studies in cognitive and perceptual development. Cambridge University Press.

Quartz, SR & Sejnowski, TJ (1997). The neural basis of cognitive development: a constructivist approach.